# Infrastructure Strategy

*Virtual and Physical Infrastructure Approach*

*Date: February 28, 2014*

*Author: Kurt Bosch*

**Contents**

## Tables

## Figures

# 1 Executive Overview

The State of Connecticut has engaged SHI Professional Services to develop an overall Data Center Infrastructure Strategy for the State's Department of Administrative Services/Bureau of Enterprise Services (DAS/BEST). This document focuses on service strategy, infrastructure design and implementation as follows:

- Services Strategy – This section will discuss strategies focused on hosting, Infrastructure (IaaS) and Platform (PaaS) service strategies
- Design of Infrastructure – The goal for the design is to provide an infrastructure that meets all of DAS/BEST's major requirements and most, if not all, of their minor ones.
    - Overall Infrastructure (Physical and Virtual) - The overall design of the new data center computing infrastructure can be broken down in two ways. The first is within each type of infrastructure – physical, virtual x86, or virtual z/VM, and the second is by each infrastructure area – network, server, storage, backup, disaster recovery, and security. The first approach provides an overall design for each type of infrastructure, for both virtual and physical. The second approach provides more specific detail regarding how a major component of the infrastructure would be designed to meet the major requirements.
    - Network - The network design is limited to the Ethernet switching infrastructure as it applies to all three types of infrastructure (virtual x86, virtual z/VM and physical) with the focus on virtual x86 infrastructures.
    - Server - The server design is specifically for the design of the x86 physical servers supporting the virtual x86 infrastructure.
    - Storage - The storage design has three major areas. These are the storage design for the virtual (x86), virtual (z/VM) and physical (x86) infrastructure.
    - Backup - The backup design is intended to provide a complete end-to-end solution for backing up all of the environments that currently exist at the DAS/BEST site.
    - Disaster Recovery - The disaster recovery design is intended to take advantage of an active/active data center model and provides a mechanism for agency customers to failover/recover their applications at a secondary site.
    - Security - The security controls that DAS/BEST manages for its' customers is identified in this section in 15 key areas.
- Implementation Strategy - The implementation of this strategy discusses four key areas. Recommendations are provided to achieve efficiencies through infrastructure, charge back, automation, and configuration management.

The methods used to arrive at this strategy included initial discovery sessions with each infrastructure and application team regarding the current network, storage, and server architecture as well as current application architectures and support process. Additional data collected prior to the start of the project were reviewed and knowledge gained incorporated into the overall strategy and roadmap.

Although the State of Connecticut agency DAS/BEST has virtualized a large portion of its physical workloads, this was accomplished using vendor products already in use within the agency. The agency has an opportunity to implement a virtual infrastructure that will allow almost any workload to be virtualized, while at the same time pursuing a completely integrated approach in the design bringing it closer to the ideal of converged infrastructure. Finally, the agency should pursue the use of automation and orchestration to further reduce the need to continually engage in the manual configuration, modification and management of the infrastructure.

This strategy recommends that the State of Connecticut simplify its infrastructure by implementing products specifically chosen to integrate tightly together and require the least amount of effort to manage and maintain. In addition, the agency should automate all of the basic, repetitive tasks currently done by internal staff.

The implementation recommendation is:

- Deploy an infrastructure designed for tight integration specifically to support virtual workloads with rack mount servers, hybrid storage (low cost, high performance) and a network switch able to integrate with the infrastructure.
- Isolate each unique platform infrastructure to limit overall support effort and technical dependencies with other operating environments, further reducing downtime, the need for critical updates and overall staff effort.
- Implement a billing system to show overall infrastructure costs as used by their customers. This has the side benefit of providing documentation to the state government that shows lower costs over time and actual usage by customers regardless of transfer payments.
- Automate the most repetitive tasks required for maintaining and configuring the infrastructure and subsequently automating those remaining tasks that represent the largest amount effort. This should be an iterative process until the cost to automate the remaining tasks exceeds the current cost to complete them manually.
- Use configuration management software to automate the installation and configuration of the supporting application software used to create customer applications within the infrastructure to reduce the effort required by the staff engaged in the manual configuration of those components.

The State will realize savings over the course of the next five years in infrastructure costs by undertaking the implementation of a tightly integrated converged infrastructure, particularly with greater emphasis on virtualization to provide a more flexible infrastructure with greater utilization of computing resources.

Following the recommendations discussed in this strategy document will yield the primary benefit of reduction in effort required by the staff to engage in repetitive tasks to manage the overall computing environment. This has a secondary benefit to allow focus and pursue activities that would lead to potential innovation and further cost reduction. Additional benefits are an infrastructure capable of providing more services to their customers at an overall lower cost.

# 2 Services Strategy

The services presented here are:

- Hosting/Co-location
- Infrastructure Platform (can be either virtual and physical)
- Application Platform (Application development environment)
- Backup/Recovery
- Disaster Recovery
- Messaging
- Collaboration
- Voice/Telecom

TCO Options associated with each strategy are located in the file attachment in Appendix A.

Table 1 below represents the overall infrastructure model. The third column lists the service offering and its relationship to the service layers and infrastructure model.

<table>
<tr><th colspan="3">Infrastructure Model</th><th>Service Layer</th><th>Service Offering[1]</th></tr>
<tr><td>Messaging</td><td>Collaboration</td><td>Telecom/VOIP</td><td>[Collaboration]</td><td>Software Services</td></tr>
<tr><td colspan="3" rowspan="3">Application Platform<br>[IIS/.NET/SQL] [Server Java/MQ-Series/DB2]</td><td>[Data Services]</td><td rowspan="3">Platform Services</td></tr>
<tr><td>[Application Hosting]</td></tr>
<tr><td>[Architecture Services]</td></tr>
<tr><td colspan="3">Disaster Recovery</td><td>[Disaster Recovery]</td><td>Disaster Recovery Services</td></tr>
<tr><td colspan="3">Backup Recovery</td><td>[Backup/Recovery]</td><td>Backup Services</td></tr>
<tr><td colspan="3" rowspan="3">Infrastructure Platform<br>[Linux, Windows, zOS] [VMware, z/VM]<br>[Servers, Storage, Network]</td><td>[Platform – OS]</td><td rowspan="3">Infrastructure Services</td></tr>
<tr><td>[Platform – Virtual]</td></tr>
<tr><td>[Hardware]</td></tr>
<tr><td colspan="3">Hosting/Colocation<br>[Power, Cooling, Floor]</td><td>[Facilities/Data Center]</td><td>Hosting Services</td></tr>
</table>

**Table 1 – Services Model**

These services also represent a continuum of services, since each service is built from the service below. This means that successive services at a higher layer depend on the service(s) below it. For example, Infrastructure services must have a data center (hosting service) in which to operate and platform and software services must have an infrastructure on which to run. (See diagram Figure 1 below).

The last two services offered, backup and DR are necessary for the protection of customer's data. These two services are also related to each other, since both provide data protection for a customer's data or application though in different ways. Backup services provide protection against lost or corrupted data and DR provides protection with a means to recover an application to a second remote site.

[1] The industry currently uses the phase: As A Service or (X)aaS to describe these services as it applies to cloud computing. This document will avoid this convoluted designation and simply discuss a particular service, such as Infrastructure, Platform or Hosting.

This list of services represents all of the services that DAS/BEST should offer to its customers, with, perhaps, one slight modification. The infrastructure services described here refer to virtual infrastructure, not physical infrastructure. However, some customers will still insist on running physical infrastructure for their applications, but managed or supported by DAS/BEST within their data center. Traditionally this is referred to as a managed service, in this case for physical infrastructure. Although it is likely that customers will request DAS/BEST provide this service, we recommend that they should only provide it when all options for virtualization have been examined and specific technical reasons to avoid virtualization are presented by the customer. The primary goal is to virtualize as much of the infrastructure as possible as quickly as possible and to offer only platform and software (Collaboration) services on a virtualized infrastructure platform.

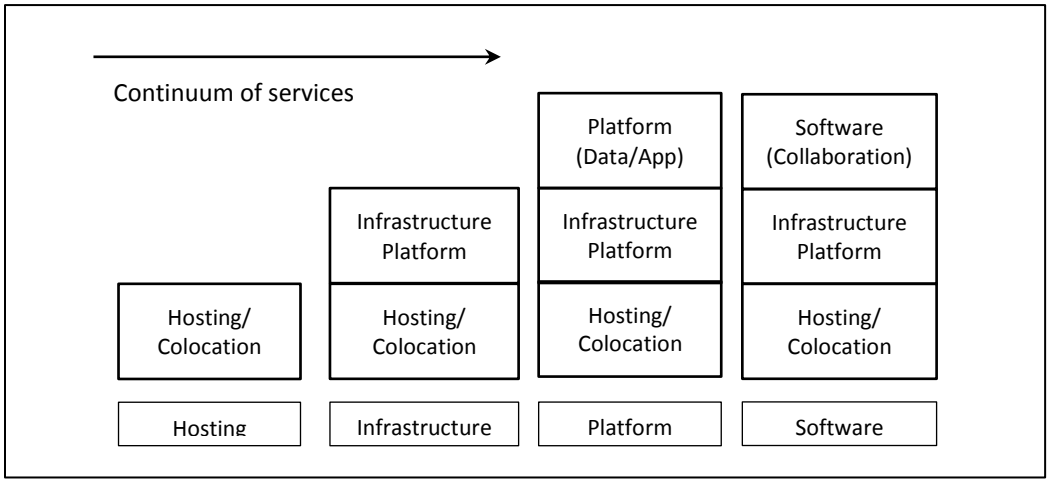


**Figure 1 - Services Continuum**

The two services Backup and DR apply to only two of these services – infrastructure and hosting. A customer that uses hosting services might also decide to take advantage of backup services, however DR services would be unavailable since that would require the customer use DAS/BEST infrastructure. A customer that uses (virtualized) infrastructure services would also be able to use backup and DR services. A customer using physical infrastructure would have limited DR functionality, though probably complete backup services. Software (Collaboration) services (Exchange, SharePoint, and Unified Communications) may or may not have these features, backup and DR, built into their design based on customer requirements.

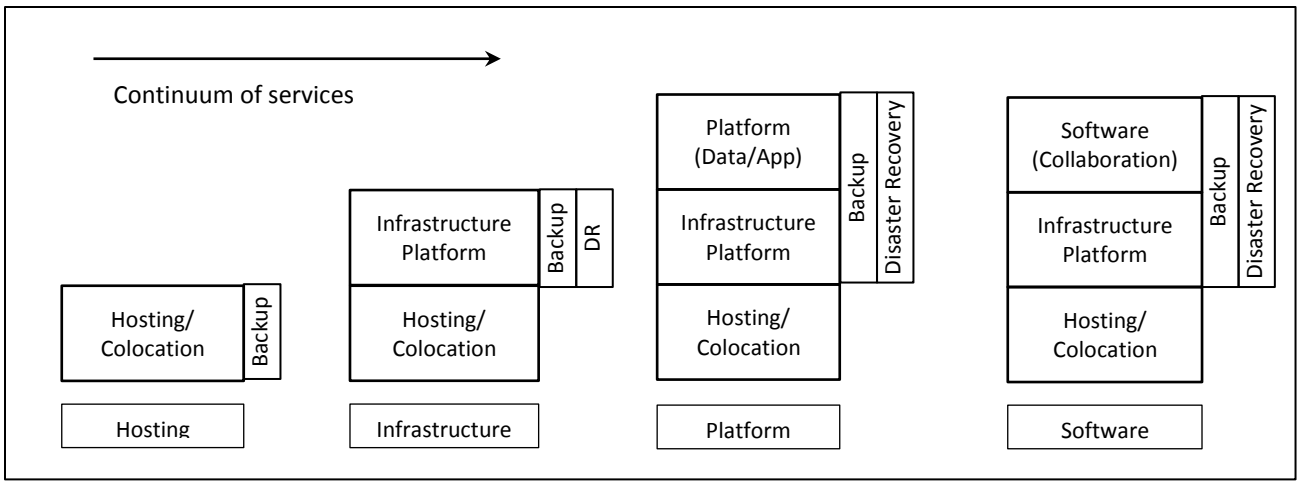


**Figure 2 - Backup/DR Services**

Platform services would have these options available to those customers using DAS/BEST infrastructure. The services continuum diagram is recreated with the DR and backup services shown as they relate to each service offering in Figure 2 above. In sections that follow, we discuss specific service offerings.

Each section describes the specific service strategy or offering, the limits of the offering, how the customer might be charged and on what basis (daily, weekly, monthly) and a brief analysis of how costs could be determined.

## 2.1 Hosting Services (Co-location)

Hosting services refers to the provisioning of floor space and power and cooling to host and support a rack (or racks) of equipment. Many different providers, notably SAVVIS, Digital Realty, Terremark, Rackspace and Equinix, offer this type of service. For DAS/BEST, given the data center being refurbished in Groton, CT, this is a straightforward co-location service offered to their agency customers.

The cost for this service is straightforward; the only specific unknown is the amount of power used by a particular customer. However, since each rack would require some number of 30 amps power cables to power the PDU's in the rack this could be estimated. Other possibilities include power metering within the rack using Metered PDU's from APC. With the use of Data Center Infrastructure Management (DCIM) software available from several vendors, power usage could be tracked for all of their customers. The only other unknown is a rough estimate for the cooling costs associated with a specific customer's equipment. Using a standard/average value for the PUE[2] based on total power consumed in running the datacenter averaged each month, would provide the cooling and overhead energy costs within the data center.

To provide this type of service, specific limits and an approach to determining costs should be set. In this case, given that the customer might require resources to power a single or multiple racks, limitations or

[2] PUE – Power usage effectiveness. PUE is the ratio of total amount of energy used by a computer data center facility to the energy delivered to computing equipment.

constraints would be the amount of floor space assigned to a rack, the type of rack used, the amount of power supplied and physical specifications for the rack itself including height, depth and width. A diagram showing a possible floor layout is shown in Figure 3.

In this particular case, floor space assigned to each rack is eight tiles with airflow from front to back. Each rack would have two 30 amp power circuits (not shown but routed through the rear cutout) to support equipment within the rack such as servers, storage and network equipment. Other constraints placed on the customer would be the service level agreement to provide uninterrupted power, provision the rack within a certain time to stand up the rack on the floor and provide power. Each rack would be provided with power distribution within the rack and any other ancillary equipment required for the rack to be properly integrated into the data center such as patch panels or mounting hardware.

The cost would be based on the floor space consumed, in our case 24 Sq. Ft., and the power provided at 0.25 Amps/Sq. Ft. with some metering provided within the rack to determine power costs over the course of a given period – either weekly or monthly. Finally, DAS/BEST would provide network access to the agency's internal network. Table 2 below provides an example of the constraints placed on the customer, service level agreement parameters provided by DAS/BEST and costs associated with the service.

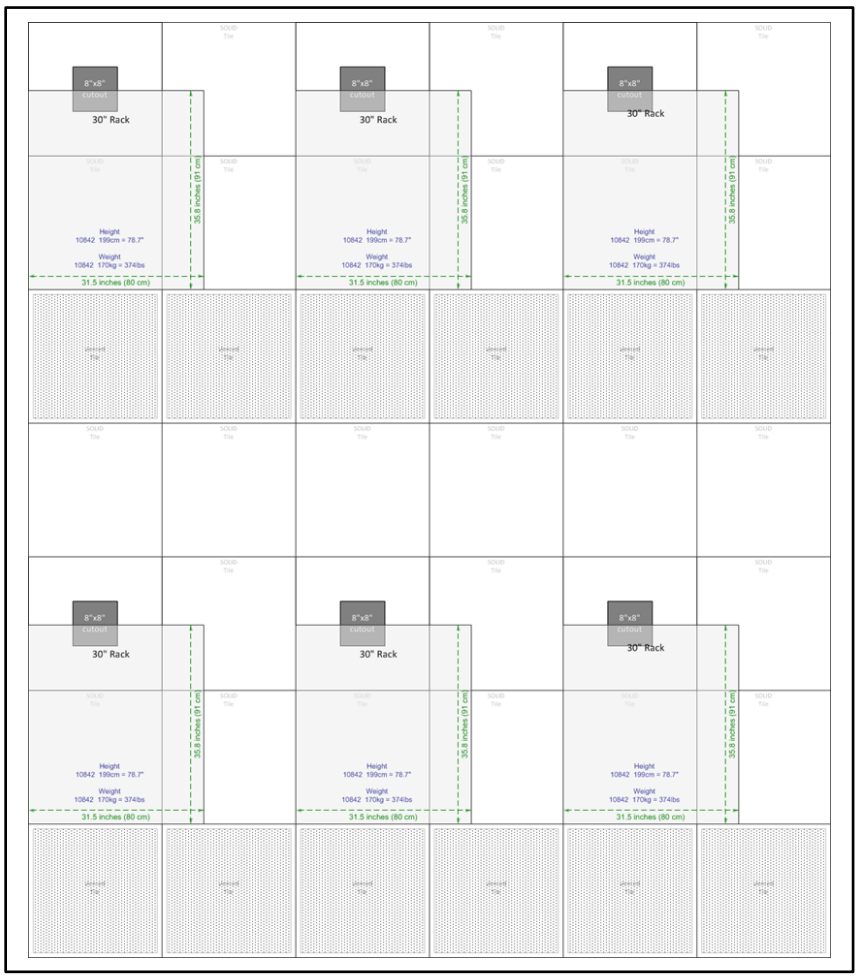


**Figure 3 – Rack Floor Layout**

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Agreement* | Time to Provision | Days or weeks, depending on current processes | This effort would be reduced over time |
| | Uptime Power | 99.9999% or 0.5 Minutes without power each year as an example | Ultimately, this should be close to zero downtime with the use of UPS systems |
| | Network Uptime | 99.99% uptime or 52.5 minutes of downtime each year as an example | This might be broken down by quarter or accumulated over the course of a year A reasonable value acceptable to customers would need to be used |
| *Service Costs* | Floor Cost | Cost for a single rack based on floor tile usage (roughly 6-8 tiles) | The prorated depreciated cost of the data center for a square foot of floor space – each rack would require roughly 24 sq. ft. |
| | Power Cost | Either using a straightforward average based on provisioned power or actual usage based on measurements gathered from metered PDU's | Determining the actual cost may prove initially to be very difficult without the use of DCIM software |
| | Network Cost | Cost of using and/or accessing network services | This might be based on port cost or throughput cost |
| *Service Constraints* | Space | The rack units (RU) available to the customer to place equipment – from 32U to 42U | This would be the rack dimensions provide to the agency customer (width, height) |
| | Power | Roughly limited to 30-40 amps each rack | This would be the maximum power provided to each rack based on total power available to data center. |
| | Network | The number of ports provided to agency, and type of switch and/or model used | This would be the limit on the number of ports allocated from main core switches to provision network access to agency rack. Other limits would be the vendor and model switch that could be used to connect to infrastructure |

Table 2 - Hosting Service (SLA, Costs, Constraints)

## 2.2 Infrastructure Services (IaaS)

Infrastructure Services in this section will mean the infrastructure available for use by agency customers of DAS/BEST. In this model, there are several infrastructures in use. The first is virtual infrastructure composed of VMware, rack mount servers, modular hybrid storage and network. The second is physical infrastructure, purchased and supported on behalf of the agency customers. The last is virtual infrastructure composed of IBM's z/VM hyper-visor, mainframe and storage.

### 2.2.1 Service Costs

As a strategy DAS/BEST will base the cost for the service of virtual infrastructure on the overall cost of the hardware and software necessary to support the environment and the labor costs to maintain and operate the infrastructure along with the cost of overhead. The service would allow for the use of a single or multiple virtual operating environments (virtual guests) of various pre-determined sizes. For example, there might be three sizes of x86 VM's available, small, medium and large. The small size would have a single vCPU and 4 GB of memory assigned. The medium size might be twice as large and the large size would be twice as large as the medium. The best approach to determining these sizes for

this service would be the review of the current sizes in use by agency customers. Most of the current VM's will be within certain boundaries, from these, choices can be made that closely approximate what would be requested.

Alternatively, the cost for a preset amount of vCPU and memory resources could be calculated and then each virtual guest would use some multiple of this amount. For example, if the base unit is 1 vCPU and 2 GB memory at a specified cost of $30/month, then if the agency customer requests a virtual guest that is four times this size (4 vCPU and 8 GB RAM), their cost would be four times the base cost or $120/month. Whichever method is chosen, a review of the current environment is necessary to gauge what sizes of virtual guests agency customers might request.

The other aspect of the cost for the infrastructure is the floor space and power used, since the infrastructure occupies space within the data center. However, this has already been determined as part of the hosting service offering as outlined above. Since this cost is based on a single rack within the data center and the virtual infrastructure occupies space within the rack, then it is only necessary to take the hosting cost for a single rack and spread that cost across all of the resources within the virtual infrastructure. As an example, consider the case where sixteen 2U rack servers are installed in a single rack with network switches installed at the top of the rack. Since this environment is "hosted" within the data center, the cost for providing power, floor space and cooling are already known. If a preset base virtual machine is used as the initial cost, for example 1 vCPU and 2 GB of memory, then this particular rack can host 2,048 VM's.

If the cost of hosting were $2050/rack/year, then each VM would absorb a cost of $1 in addition to hardware and labor costs to purchase, maintain and manage the virtual infrastructure. Assuming the virtual infrastructure within a particular rack (purchased as a unit) were approximately $1.8M[3], then the base cost for the preset VM would be $300/year (assuming the overall costs were prorated over a 3 year support period and support for 3 years was purchase initially). If the labor costs are added in based on some initial assumptions about staffing, such as one full time administrator for each 1024 VM's and a labor cost of $123,000, then the labor costs would be, roughly, $10/VM/month.

[3] This would include all software licenses for VMware as well.

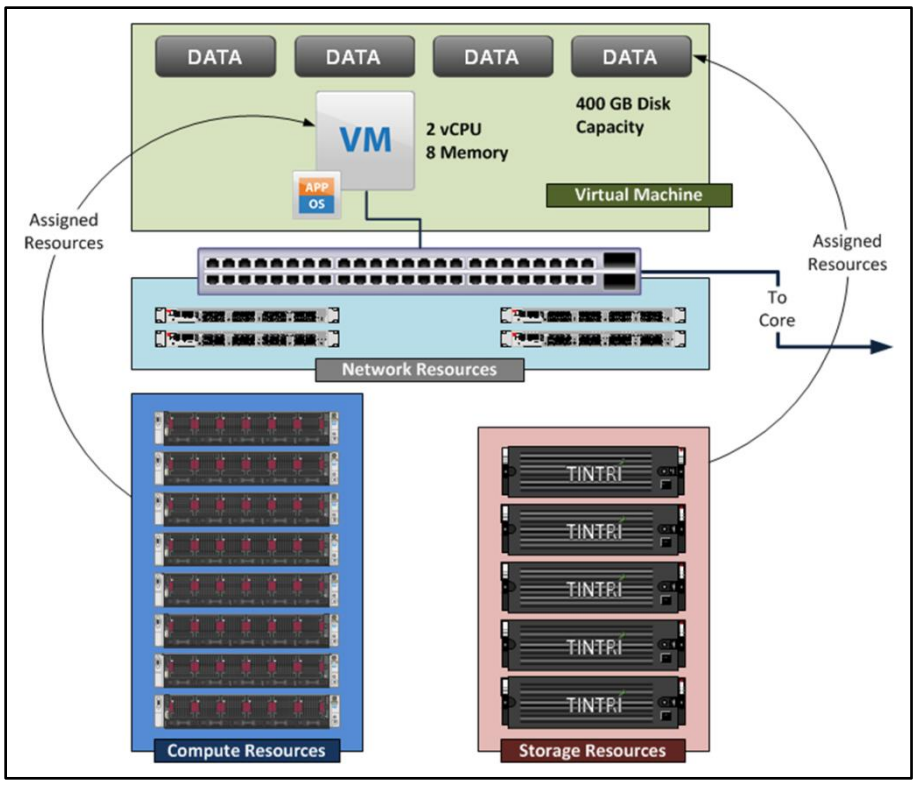


**Figure 4 - Virtual Machine Assigned Resources**

Now adding all costs so far, $1/year for hosting, $300/year for the virtual machine hardware and software (base amount), $120/year for staff or $321/year[4] for the cost to host a single small VM. Monthly this would be $27.

Other costs, such as staff overhead or management and network costs have been omitted. However, taking these additional costs, and spreading them out over the infrastructure just as outlined above, will determine the final cost based on the initial VM. Storage costs are treated separately by determining the overall cost to manage, maintain, purchase and support and then determine the cost for a specific fixed amount, such as a 100 GB. For example if all of the costs (hardware, software, and labor) associated with storage is $2.38M for 330 TB's of storage over 3 years, then a 100 GB disk would cost ~ $240/year or $20/month. Figure 4 above depicts the construction of a virtual machine constructed from compute, storage and network resources.

For a customer requesting a particular virtual machine (VM) with four vCPU, 8 GB memory, and 500 GB storage, the monthly and yearly costs are respectively $208 and $2,496 for this VM. The cost of network could be added in based on a prorated amount associated with the number of ports or bandwidth committed to the virtual machine. A similar exercise can be conducted for the Mainframe virtual infrastructure as well; yielding a final cost for a base preset VM on which a customer could base their order. Beyond the cost of this service, there is still the question of what service agreements and constraints on the services should be offered.

[4] Amazon charges approximately $500/year for a single small instance using 1 vCPU and 1.7 GB memory w/160 GB Storage.

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Costs* | Host/Guest | Cost for a single base VM or several prices for varying sizes | Based on the cost of the hardware, software and staff |
| | Network | Cost for a specific amount of bandwidth | Based on prorated cost for a single customer based on a shared model |
| | Server | Cost for a particular server model from a limited selection | Based on a shared staffing model and the use of a particular vendor hardware and limited to specific options |
| | Storage | Cost for a specific amount of storage in some base unit (100 GB as an example). | Based on a shared staffing model, includes the hardware and software costs with the price based on an incremental unit |

Table 3 - Infrastructure Service Costs

### 2.2.2 Service Level Agreement

The service level agreements also need to be determined for virtual infrastructure, though in most cases the agreements, baring special cases, should be straightforward. In general, the customers are likely to care about several different aspects of the service provided. The first is the amount of time to provision a single VM with the agreement that virtual infrastructure will be stood up within a few hours of the initial request. This type of request can be automated using tools available within most virtual infrastructures so that requests for complex environments comprised of many VM's of varying sizes can be deployed quickly with environments composed of certain software components already installed and ready as templates.

The second service agreement will be the uptime guarantee for any virtual machine running within the infrastructure. Service level guarantees vary among companies providing virtual compute services from as low as 99.95% (4.3 hours of downtime/year) to as high as 99.999% (5 minutes of downtime/year). Given the nature of virtual infrastructure and experience with current customer commitments, a reasonable uptime guarantee can and should be determined. Finally, there might be service level agreements to address how quickly changes might be made for particular aspects of their virtual infrastructure, how quickly certain types of potential issues will be resolved,

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Agreement* | Time to Provision | Within a few hours or a day depending on current internal processes | This effort would be reduced over time |
| | Network Uptime | 99.99% uptime or 52.5 minutes of downtime each year as an example | This might be broken down by quarter or accumulated over the course of a year A reasonable value acceptable to customers would need to be used |
| | Host Uptime | 99.99% uptime or 52.5 minutes of downtime each year as an example | Same as above |

Table 4 - Infrastructure Service (SLA)

### 2.2.3 Service Constraints

Service constraints for infrastructure services are based, primarily, on the size and type of virtual guests that can be provisioned. VMware has specific limits on the maximum size of virtual guests that can be created – 32 vCPU's and 1 TB of memory. Additional constraints may be placed on the virtual guests

based on network bandwidth and topology. However, since this is virtual infrastructure there are only a few other constraints. Additionally, some constraints may apply to storage and network performance.

Physical infrastructure may have several more constraints, including speed and number of network connections and limits on specific vendor x86 hardware models that can be used within the environment.

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Constraints* | Host/Guest | Maximum size 32 vCPU's, 1 TB memory | Limit here is the maximum size of the VM and how the VM is initially sized. |
| | Network Bandwidth | Maximum bandwidth available 1/20 of overall bandwidth available to host. (approx. 1 Gbps) | Limit is the portion of the bandwidth available to an individual VM. QoS limits may also be used. |
| | Physical Server | Specific sizes, models and features | Constrained to use potential models and options as determined by DAS/BEST. Based on current models in use. Will depend on future plans. |
| | Storage | Storage increment might be 100 GB with additional increments of 100 GB to a maximum of 2 TB's. | The storage would be constrained to a specific increment or size. Customer can increase using a specific increment with a potential maximum limit enforced. |

Table 5 - Infrastructure Services (Constraints)

## 2.3 Platform Services (PaaS)

In this section, platform services refers to the actual underlying software components used to "build" applications such as the various database technologies (Oracle 11g/12c, IBM DB2, and MS SQL Server), and others (MS IIS, IBM Websphere, and Oracle Weblogic). In the proposed infrastructure, each of these would run on a specific subset of the overall infrastructure. For example, the MS components would obviously run on the virtual x86 infrastructure, with some or all of the IBM and Oracle components also virtualized on x86 hardware.

### 2.3.1 Service Costs

Service costs for platform services are priced using a similar approach as that outlined for infrastructure services, though in this case, there are no hardware costs to include[5]. For the software, it is a simple matter to assign the license costs for the solution. The difficult part is determining the cost for the final service offered to the customer since this is based on both the actual software and the staff costs. The simplest approach would be to assign a portion of the staff cost to the support of each database instance currently running in the data center. This could be done based on either the size of the overall database (storage), the number of users the database supports, or the size of the database environment (CPU).

As an example, assuming an initial cost for the staff of $3.0M a year and 300 instances to maintain and support, the cost for each database as a portion of the staff costs might be calculated as simply as

[5] This is not strictly true, though it depends on the approach DAS/BEST chooses to take. If the infrastructure for the Oracle platform is done separately from the software licenses, then the same approach can be taken, however, it may be a bundled solution and it may not be possible to break out the hardware and software.

$10K/instance/year or $833/instance/month. If instead, each database instance consume a certain amount of storage and there are 300 TB's under management, then each terabyte of storage for a particular database would incur a charge of $833/TB data/month. Finally, using CPU counts, which more closely align to license costs, assuming 300 CPU's assigned to database instances would yield the same number on an individual CPU basis. Further granularity can be achieved by pricing each type of database separately, though determining the portion of the staff costs to assign may prove difficult. Ultimately, the costs should be assigned to each type to best identify the least expensive platform.

Assuming this is MS SQL Server, then the cost based on whatever approach above DAS/BEST chooses, would be significantly less than the others. Ultimately, this serves as an enticement to customers to move to the lowest cost platform assuming it meets the requirements of the proposed application. Over time, the economics would help DAS/BEST move to an overall lower cost platform for its customers. Table 6 provides an example breakdown that makes this point, albeit artificially.

| Database Vendor (Type) | Storage (TB) | CPU Usage | Staff Costs | License Cost (CPU) | Customer Costs /month | Total Monthly Cost[6] |
|---|---|---|---|---|---|---|
| *Oracle* | 50 TB | 20 CPU's | $1M | $3M | $12,500/CPU License<br>$1,667/TB Labor<br>$1,667/CPU Labor | $14,167 |
| *IBM DB2* | 50 TB | 30 CPU's | $1M | $2M | $5,556/CPU License<br>$1,667/TB Labor<br>$2,778/CPU Labor | $8,334 |
| *MS SQL Server* | 300 TB | 150 CPU's | $2.4M | $600K | $334/CPU License<br>$667/TB Labor<br>$1,334/CPU Labor | $1,667 |

**Table 6 - Database Instance Costs**

However, this is only the costs for the database services, a similar analysis can be conducted for the other components of a particular software platform, with costs for licensing assigned to a specific instance and the cost for staff spread across the various environments using a particular approach as that chosen for the database services. This still, however, does not include the costs of the underlying layers that support the platform such as hosting and infrastructure. Since we already examined those costs and determined a cost for a base VM, we can simply use that cost for each of our virtual guests supporting a particular instance and then combine them for a particular platform.

For example, suppose a customer requires a platform composed of IIS servers, additional servers to handle connections from IIS to a SQL Server with a single database instance all running within the virtual infrastructure environment as we described above. We could, with specific information about the various sizes and numbers of the virtual guests required, compute a cost to provision and support that environment for a particular agency customer. Table 7 below outlines these costs and lists the total based on some assumed costs for the platform software components.

| Platform Component | Size (VM) | Qty | Base Cost (Monthly) | Total Cost (Monthly) |
|---|---|---|---|---|
| IIS Web Server | 2 vCPU<br>8 GB RAM<br>100 GB Disk | 4 | $54 VM<br>$20 Storage<br>$667 IIS | $2,964 |
| Other | 2 vCPU | 2 | $54 VM | $776 |

[6] The total uses the larger of either the storage capacity labor cost or the CPU labor cost to determine the final monthly cost.

| | | | | |
|---|---|---|---|---|
| | 8 GB RAM | | $334 Other | |
| Database (SQL) | 8 vCPU<br>16 GB RAM<br>400 GB Disk | 1 | $216 VM<br>$80 Storage<br>$13,336 DB | $13,632 |
| | | | ***Total*** | ***$17,372/month*** |

**Table 7 - Sample Customer Platform Costing**

The various components included in the services costs are given in Table 8. Other refers to any other valid component used within the platform to provide a base on which the customer can build an application. The cost for a single instance is based on the software license cost and the staff or labor costs spread across all of the instances.

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Costs* | Database | Cost for a single instance of a preset size or several prices for varying sizes | Based on the cost of the software and staff (hardware costs rolled up as part of VM cost) |
| | Web Server | Cost for a single instance | Same as above |
| | Other | Cost for a single instance (includes all licensing) | Same as above |

**Table 8 - Platform Service Costs**

### 2.3.2 Service Level Agreements

Service agreements for platform services are even simpler than those services discussed above. Only two specific service agreements are required, one for provisioning and one for uptime. The provisioning should occur within a specific period, perhaps as short as a few minutes or hours. The service level for uptime would be stated in a similar fashion as that for infrastructure. For a complex platform an uptime SLA ranging from 99.9% to 99.99% might be used.

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Agreement* | Time to Provision | Within a few hours or a day depending on current internal processes | This effort would be reduced over time |
| | Platform Component Uptime | 99.99% uptime or 52.5 minutes of downtime each year | Same as above |

**Table 9 – Platform Service (SLA)**

It is likely that, over time, other aspects of the service level agreement will be required. These might include the time to upgrade platform software components and time to apply patch updates.

### 2.3.3 Service Constraints

The service constraints for platform services are much different from those for infrastructure or hosting. Those services are constrained by their physical aspects. Hosting is the aspect of actual power that can be delivered; for infrastructure, the available resources that can be provided to a virtual or physical host.

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| *Service Constraints* | Database | Oracle 12c/11g, IBM DB2, MS SQL Server 2012 | Limited to a short list of major databases. This may be reduced down to two over time. |
| | Web Server | MS IIS, Oracle Weblogic, IBM Websphere | Limited as above. This list may be reduced over time. |

| Service Area | Item | Comments | Notes |
|---|---|---|---|
| | Other | JAVA, .NET | Same as above. |

Table 10 - Platform Services (Constraints)

For software-based services, the obvious constraints are what specific products should be made available. Primarily because, given a limited staff, there is only so much software that any group of professionals can support. The list of products supported may be reduced over time; however, it may also change or, as efficiencies are achieved using automation, be increased.

## 2.4 Infrastructure Cost Model

The cost model for the infrastructure is based on the various services DAS/BEST will offer to its agency customers. At this point only hosting and infrastructure and platform services are represented, several additional services are identified earlier in the strategy. These are Disaster Recovery services, Backup/Recovery services and Software services. In this section, a breakdown the cost structure for the organization is presented and how we suggest it should be organized to assign costs to the various services offered.

Each service is composed of some hardware and/or software components and staff that support that particular service. In the case of virtual infrastructure, it is both the hardware, virtual hyper-visor, operating environments (Windows or Linux) and the staff that handles daily support of the physical infrastructure. This is similar to the items that make up the costs for the physical server infrastructure.

Table 11 shows the hardware and software cost items associated with each type of service. The final column lists the cost metric or basis for charging for the service. For storage, this is the cost for each 100 GB capacity requested by a customer. For network resources, associated with each of these other services, the cost for a 1 GbE link/bandwidth is added to the underlying infrastructure cost.

| Service | Hardware/Other | Software Costs | Staff | Cost Metric |
|---|---|---|---|---|
| Hosting | Floor space<br>Power + A/C plant, UPS<br>Cabling + Racks<br>Building (Datacenter) | DCIM Software | Facilities Staff<br>NOC Staff | Rack/Floor Cost (24 Sq. Ft.) |
| x86 Virtual Infrastructure | Rack Servers<br>Storage | VMware Enterprise + | VMware Admins | Storage Cost (100 GB) |
| | | Microsoft OS<br>RedHat Linux OS | System Admins (MSFT, Linux) | |
| | Network Switches | | Network Admins | Bandwidth Cost (1 GbE) |
| z/VM Virtual Infrastructure | Mainframe z114 | z/VM Software<br>RedHat Linux OS | System Admins (MF, Linux) | VM Cost (CPU + Memory) |
| | Storage<br>FC Switches | | Storage Admins | Storage Cost (100 GB) |
| | Network Switches | | Network Admins | Bandwidth Cost (1 GbE) |
| Physical Infrastructure | Rack Servers | Microsoft OS<br>RedHat OS | System Admins | Server Cost |
| | Storage<br>FC Switches | Storage Software | Storage Admins | Storage Cost (100 GB) |

| | | | | |
|---|---|---|---|---|
| | Network Switches | | Network Admins | Bandwidth Cost (1 GbE) |
| Backup Infrastructure | Rack Servers | Backup Software | Backup Admins | Backup Storage Cost (100 GB) |
| | Storage<br>FC Switches | Storage Software | Storage Admins | |
| | Network Switches | | Network Admins | |
| Disaster Recovery | VM (recovery)<br>Storage | Storage Software<br>Replication Software | System Admins<br>Storage Admins | VM Cost<br>(1 vCPU x 4 GB)<br>Backup Storage Cost (100 GB)<br>Bandwidth Cost (1 GbE) Access |
| | Network Switches | | Network Admins | Bandwidth Cost (1 GbE) Replication |

**Table 11 - Cost Items for Infrastructure**

Cost for infrastructure and platform services are additive and include the costs for lower level services required to support the higher-level service. For example, for virtual infrastructure, the cost of hosting the environment is added to the cost of deploying the virtual infrastructure. For platform services this cost would include the cost for the virtual infrastructure (by VM), which already includes the cost for hosting the environment (hosting services).

In the case of backup services, this is the total cost to backup and storage for each 100 GB of capacity based on a predetermined retention policy as part of the backup service level agreement (SLA). The backup service cost would include all of the hardware, software and hosting (floor space, power, and cooling) costs necessary to support it.

### 2.4.1 Overhead Costs

There is still one aspect of the cost to manage the infrastructure that has not been considered; the inclusion of all overhead costs associated with the overall data center and its management. Specifically, the costs associated with staff not directly involved in daily tasks to support these services are not included, or facilities that do not directly support the infrastructure such as office space. These costs would need to be distributed among the various services offered in a way that was reasonable. Several obvious approaches exist. The first approach is to assume these costs are split equally among all of the services. The second is to assume some fractional allocation based on the amount of staff being managed.

As an example, if the staff for platform services were 50% of the total staff available, then this service would accrue 50% of the overhead in addition to all of the base costs to deliver the actual service. Table 12 shown below, breaks down an example showing how these overhead costs would accrue to each service assuming a specific percentage of total staff associated with each service. As part of the exercise, it is assumed that all facilities costs are $5M excluding the data center and all staff costs are $4M excluding all positions at the level of manager and below – these would be assigned to a specific service. These two overhead costs are added ($4M + $5M = $9M) and a percentage of the total overhead is assigned to each service based on the percentage of the actual staff required to support that service.

| Service | Staff | Accrued Overhead Cost |
|---|---|---|

| | | |
|---|---|---|
| Hosting | 5% | $450K |
| Infrastructure (x86) | 15% | $1,350,000 |
| Infrastructure (z/VM) | 5% | $450,000 |
| Infrastructure Physical | 10% | $900,000 |
| Storage Infrastructure | 5% | $450,000 |
| Platform Services | 30% | $2,700,000 |
| Backup Services | 10% | $1,350,000 |
| Disaster Recovery Services | 20% | $1,800,000 |
| **Totals:** | 100% | $9,000,000 |

**Table 12 - Sample Overhead Cost Accrual**

# 3 Design of Infrastructure

The goal for the design is to provide an infrastructure that meets all of DAS/BEST's major requirements and most, if not all, of their minor ones. First in the design are the requirements that DAS/BEST has decided must be met, both major and minor. Second in the design is the overall design for infrastructure, both physical and virtual with a heavy emphasis on virtual infrastructure. The design will provide detail for each of the major areas including network, server, storage, backup, disaster recovery and security.

## 3.1 Overall Requirements

The design of the infrastructure must follow from the overall requirements. Those requirements are that the infrastructure supports an active site at two locations, provide for a converged infrastructure model, and support a programmable model using RESTful API's. In addition, there are several other ancillary requirements, including the goal of simplicity, ease of maintenance and management, and the ability to virtualize all of the current customer's physical workloads.

### 3.1.1 Active/Active Data Centers

The requirement for active/active data centers specifically addresses the use of two data center sites to support disaster recovery. This is related to the need for a recovery site for some portion of the current customer's applications, with the requirement that some of the infrastructure at both sites be in use with the remaining idle infrastructure at either site be available for failover of critical applications from the other site. Therefore, each site maintains active applications and idle capacity for the failover of critical applications.

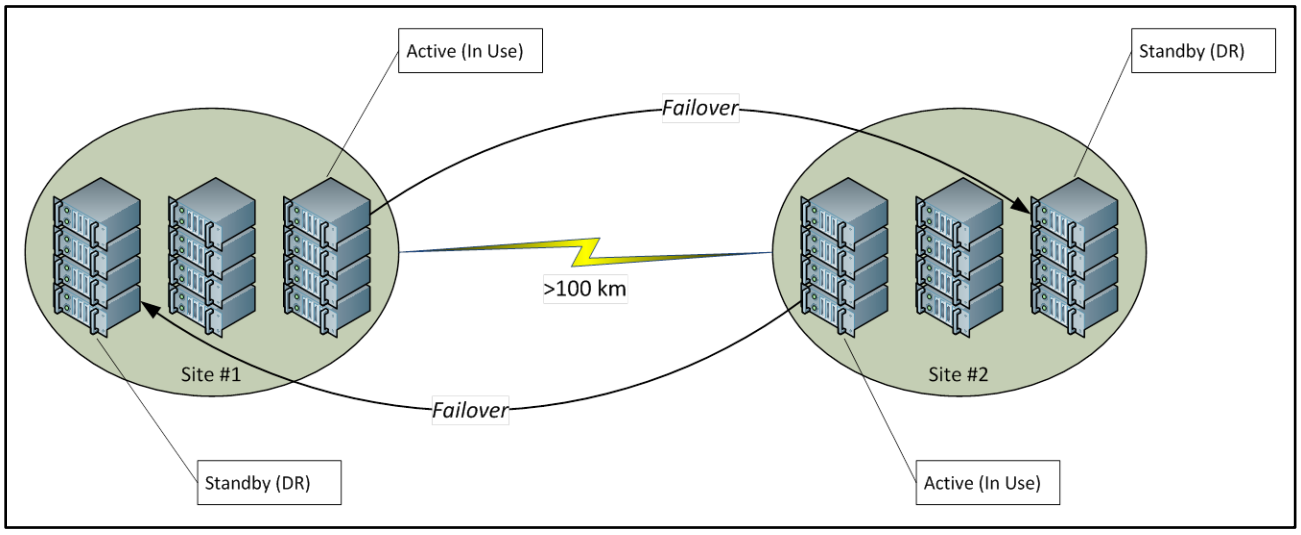


Figure 5 - Active/Active Data Center

While it is required that the design support an active/active infrastructure at two sites, there is no stated requirement that the agency customer's actually use it, though compared to current methods, the method proposed in this design should provide a much faster recovery time and closer recovery point. Figure 5 depicts how this environment would look.

The ability to actively run an application simultaneously at two sites with instantaneous failover would be prohibitively expensive and require strict software architecture guidelines. This particular approach

will not be discussed here. However, the infrastructure design as outlined below does not prohibit this approach, it would, however, require the software architecture practices be strictly defined with this in mind.

### 3.1.2 Converged Infrastructure

The definition used for converged infrastructure is provided below along with a definition for orchestration as it applies to a converged infrastructure.

> ***Converged infrastructure*** *packages multiple information technology (IT) components into a single, optimized computing solution. Components of a converged infrastructure solution include servers, data storage devices, networking equipment and software for IT infrastructure management, automation and orchestration.* ***Orchestration*** *describes the automated arrangement, coordination, and management of complex computer systems and services. (Wikipedia)*

The goal of a converged infrastructure as defined is a requirement for the overall design of the virtual infrastructure environment. The design presented below includes components for each of the hardware areas and the requisite software to provide management, automation and orchestration of the entire environment. This is provided, primarily, for the virtual x86 infrastructure and to a lesser degree, the z/VM virtual infrastructure[7].

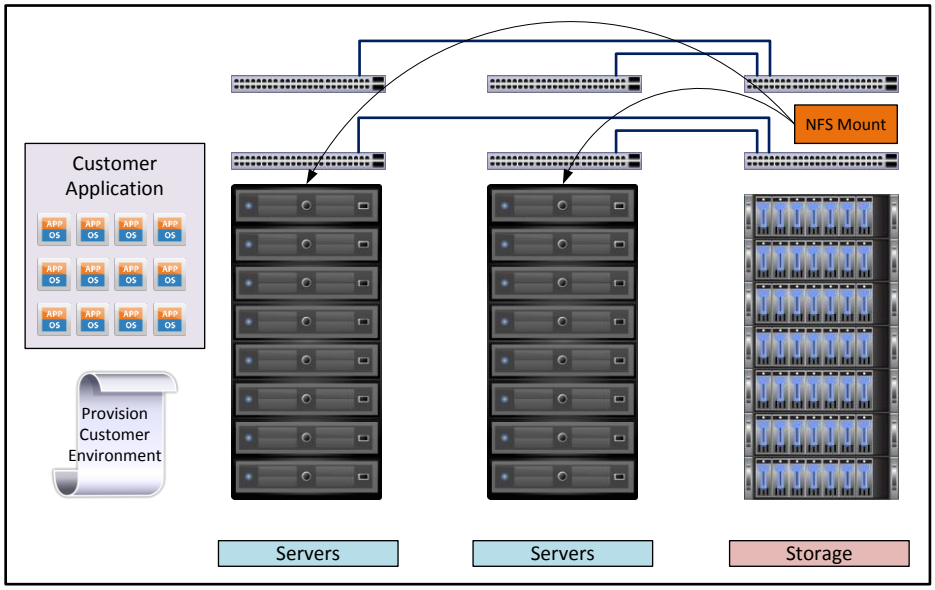


**Figure 6 - Converged Infrastructure**

Figure 6 depicts a logical view of the converged infrastructure. Ideally, the infrastructure is configured or modified based on automation to include the provision of a complete customer environment.

### 3.1.3 Programmable Components (REST)

[7] The primary difference is that z/VM environments may lack some of the necessary products to support complete automation and orchestration of the environment.

The requirement for programmable components is based on the desire to have all or almost all of the components that make up the overall design, provide an interface with a documented API using REST web services so that orchestration and automation are more straightforward to implement. This means that the network switches, storage arrays, firewalls, routers, servers, and hyper-visors that are part of the design, should provide a RESTful API or web service that can be used to configure or modify aspects of that particular component. As an example, to make changes to the network, an orchestration client would send commands over HTTP to the appropriate switch to make changes to its configuration based on predefined rules or policies defined during the provisioning of a virtual guest.

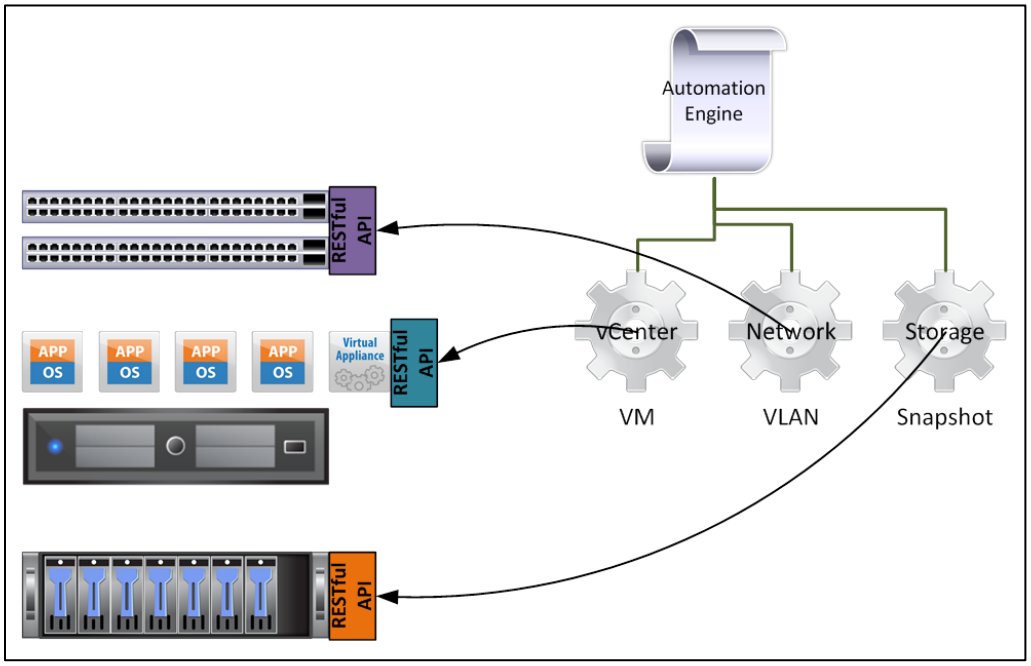


Figure 7 - RESTful API and Automation

In the proposed infrastructure, the engine that executes the automation would make execute a command using RESTful web services and the API associated with a particular component to modify either the storage, virtual infrastructure or network (see Figure 7).

### 3.1.4 Other Requirements

There are several other requirements for the infrastructures that DAS/BEST has determined are necessary. The first of these is that the proposed infrastructure must be “simple.” This means that no part of the infrastructure should have features or components that are unnecessary in supporting customer application environments.

Additional requirements are that the infrastructure be easy to maintain, manage and troubleshoot. These are required so the cost for the overall support of the infrastructure is kept to a minimum. The word “easy” is a vague description of the amount of effort necessary to accomplish certain tasks related to support and relative to the other options available. For maintenance and management, it refers to the actual number of sub-tasks that need to be completed to accomplish a major task relative to the other options under consideration. For troubleshooting, “easy” refers to the number of different components that need to be accessed to determine the root cause of the problem and the overall effort involved.

For example, if storage solution A requires 3 major tasks with a total of 15 sub-tasks on three different components and the second option, storage solution B, requires only two different components be configured with 2 major tasks and a total of 5 sub-tasks, then option B represents an easier to maintain and manage solution.

The final requirement is that all workloads currently on physical servers can be virtualized[8] on the proposed infrastructure. Typical reasons that prevent an operating environment and the associated workload from running within virtual infrastructure are the amount of resources, both CPU and memory, available within a hyper-visor host, the disk IO(s) that can be supported by the storage array and network bandwidth required by the applications. Without enough CPU, memory, disk space and performance, and network bandwidth some application workloads cannot be virtualized. The goal of this proposed infrastructure is to provide adequate resources within each of these areas (server, storage and network) to allow as many of the physical workloads as possible to be virtualized.

## 3.2 Overall Design

The overall design of the new data center computing infrastructure can be broken down in two ways. The first is within each type of infrastructure – physical, virtual x86, or virtual z/VM and the second is by each infrastructure area – network, server, storage, backup, disaster recovery, and security. The first approach provides an overall design for each type of infrastructure, for both virtual and physical. The second approach provides more specific detail regarding how a major component of the infrastructure would be designed to meet the major requirements.

The overall design of the physical infrastructure refers to the components and integration for operating environments hosted directly on physical (bare metal) hardware without the use of a hyper-visor. The overall design of the virtual infrastructure refers to the components and integration for similar environments (Microsoft Windows and RedHat Linux) running on hyper-visor based infrastructure. Hyper-visors currently used within the infrastructure are based on x86 and mainframe hardware. Each of these types of virtual infrastructure is reviewed separately.

### 3.2.1 Physical (x86) Infrastructure

The overall design of the physical infrastructure includes the servers, storage and network components that are used to support the applications. Figure 8 below shows the logical view of this physical infrastructure. Each server requires at least one connection to the Ethernet network and redundant connections to the storage network. This is a logical view of an application residing on physical servers.

Each Ethernet connection from the physical server would connect at 1 GbE to the top-of-rack switch in each cabinet. There may be more than one physical network connection depending on the application server requirements – additional connections may be required for backup traffic, inter-application communication and redundancy. Only a single top-of-rack switch is shown, typically there would be at least two switches for redundant connections from the servers. These switches would support multiple VLAN(s) specifically created for use by a particular agency customer.

Since there is a substantial investment in FC infrastructure using Brocade FC switches and most vendors support FC connection, particularly for diverse environments such as the one at DAS/BEST, there is no

[8] This does not mean that they will be virtualized, only that technically, the infrastructure could support the resource requirements for a particular workload (i.e. operating environment with an application running).

reason to replace it with any other storage network protocol. FC also benefits from having several vendors who support the protocol for storage network switches providing choice for customers. Within the physical infrastructure, these FC connections should always be redundant.

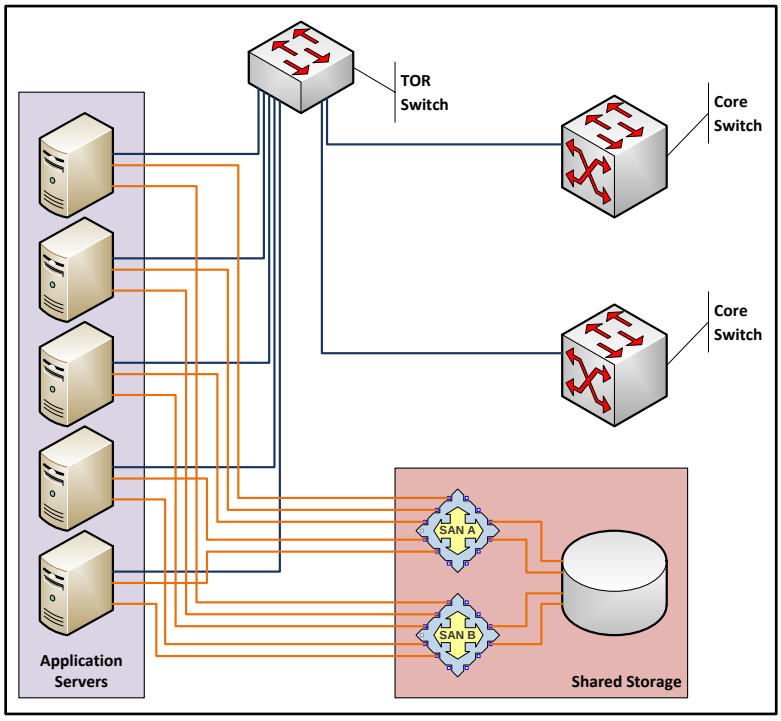


Figure 8 - Logical Physical x86 Infrastructure

Although no particular vendor is represented by the shared storage depicted, several vendors have products that will support all of the particular operating environments represented in the data center environment, a specific vendor should be chosen to replace all of the current storage that is currently nearing end of service before initial implementation of the overall physical infrastructure. Figure 9 shows the rack layout for the physical infrastructure. Placed at the top of each rack are FC and Ethernet switches to support connection to the production network and shared storage. Each server would have redundant connections to both networks with both at 1 GbE.

At the top of each rack are the access switches for Ethernet network and FC storage network connections from the servers within the rack. There are redundant connections from these switches to the core-switching infrastructure for each type of network. For the rack shown in Figure 9, several different server models are shown. Each rack would contain a random selection of servers as they are provisioned for customers. The models in use would be limited to a specific selection from certain models from one, perhaps two, vendors. The specifications of each server would be determined so that an agency customer that needed an application environment deployed would select from this list. For example, assuming the vendor is DELL; those models might be the DELL PowerEdge R520, R720 and R820. This would be the equivalent of a choice of a small, medium and large server. If HP is the vendor of choice for servers, these models might be the DL320 G8, DL3680 G8, and DL560 G8. In either case, each cabinet would have a section set aside for small, medium and large servers within each cabinet.

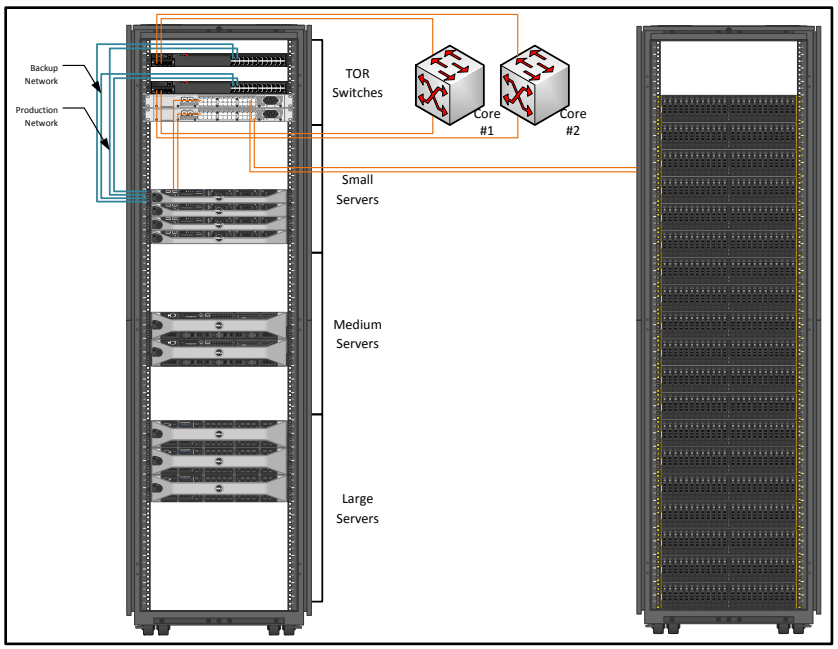


**Figure 9 - Cabinet Layout Physical Infrastructure**

Storage is dedicated to just the physical infrastructure and isolated from the other two virtual infrastructures. Shown in the diagram is an HP 3PAR storage array connected to a single server cabinet for physical infrastructure. The array shown potentially could support up to 200 Terabytes of storage in one cabinet. The physical infrastructure might include several storage arrays and many more server cabinets.

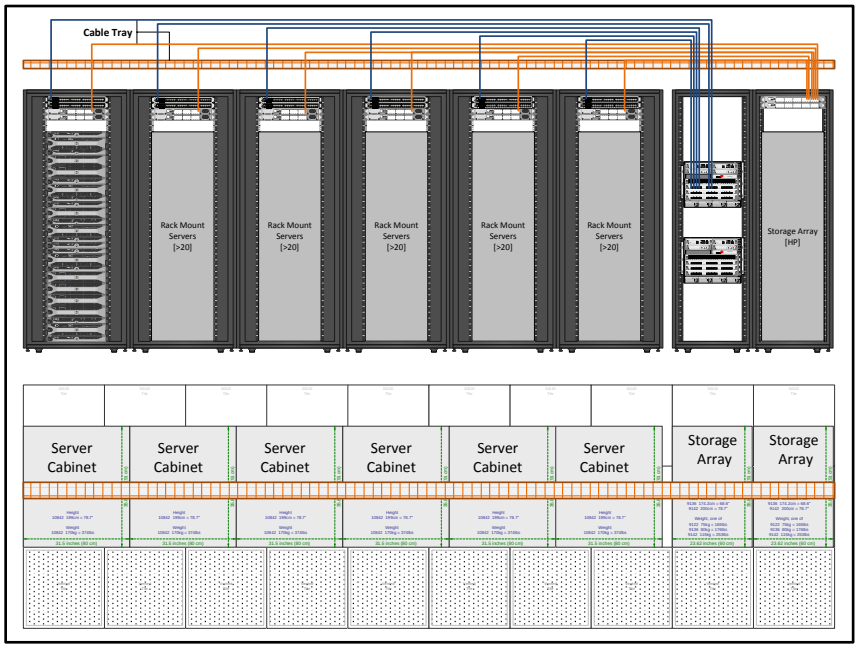


**Figure 10 - Physical Infrastructure Rack Layout**

An example rack elevation and floor layout of a possible configuration is shown in Figure 10 above. Ten cabinets could support upwards of 200 rack mount servers. Only a few additional cabinets would be

required for networking and storage. This approach would keep the physical infrastructure as simple as possible with the least number of components to manage.

### 3.2.2 Virtual (x86) Infrastructure

The virtual infrastructure is built in a fashion logically similar to the physical infrastructure with the exception that it also uses a hyper-visor to abstract the storage, network and servers and uses modular storage to support the virtual guests. An Ethernet network is used to provide access to the storage from the hyper-visor based hosts. (see Figure 11 below). For both storage and TCP/IP networks, all links are 10 GbE to support the necessary bandwidth for the many virtual guests that reside within each physical x86 server. Although only one link is represented in the logical diagram, two should be used for redundancy and additional bandwidth in each network. The storage is modular, and capacity is expanded incrementally, by adding another module as needed.

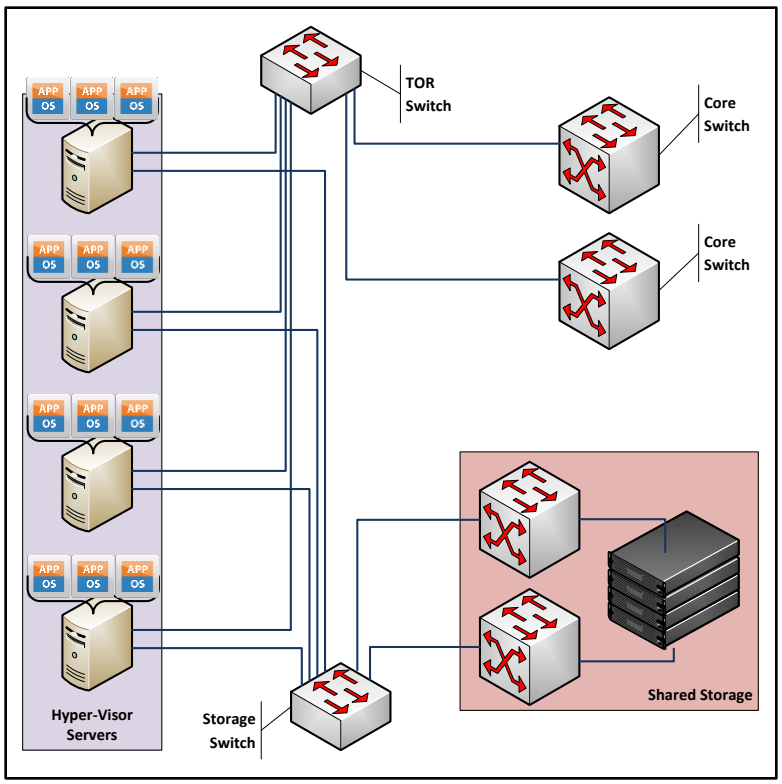


Figure 11 - Logical Virtual Infrastructure

The individual racks containing the servers are all configured in the same fashion, with 18-20 rack mount servers installed in the lower part of the rack and four top of rack (TOR) switches at the top. These switches are for production network access as well as storage network connections to the cabinet containing the modular storage. Shown in Figure 12 are a single cabinet of servers and a single cabinet of storage. Capacity is increased in each by adding more cabinets full of either servers or storage depending on the type of capacity needed. The switches shown in each cabinet have not been determined so these may be changed for slightly different models with more or less 10 GbE and 40 GbE ports depending on the final requirements for the design.

The storage network would be supported over Ethernet Switches in a full mesh fabric across six to eight cabinets. Interconnects would be 40 GbE, with local connections for servers and storage at 10 GbE.

Servers would have redundant connections to each pair of switches for production customer VLAN(s) and again for dedicated storage access. The links shown in these diagrams represent interconnect links between switches, though not all interconnects are shown. Network configuration for production would be similar to that for the physical server infrastructure. The storage network follows similar principles as for FC storage networks – redundant, low over-subscribed links, and simple topology.

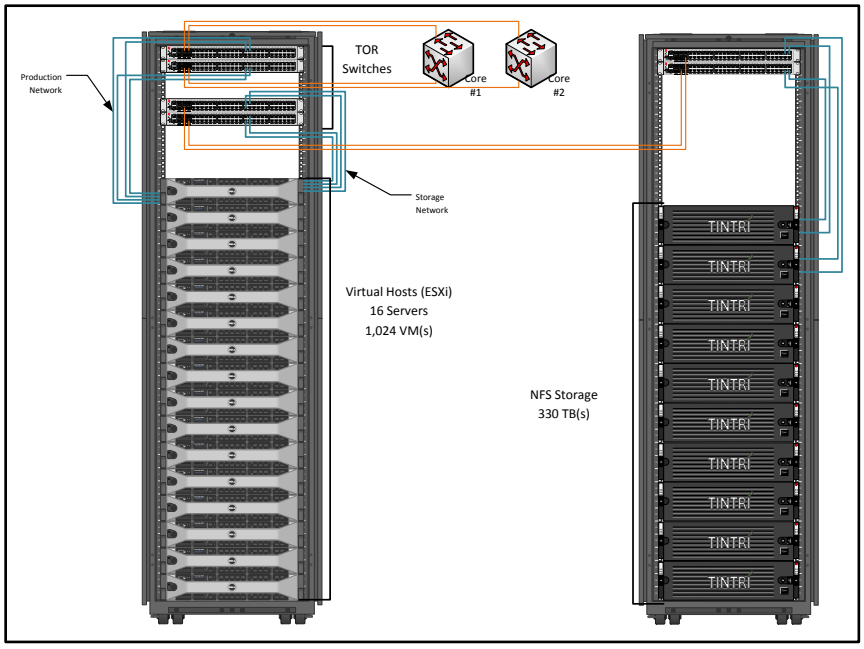


**Figure 12 - Cabinet Layout Virtual Infrastructure**

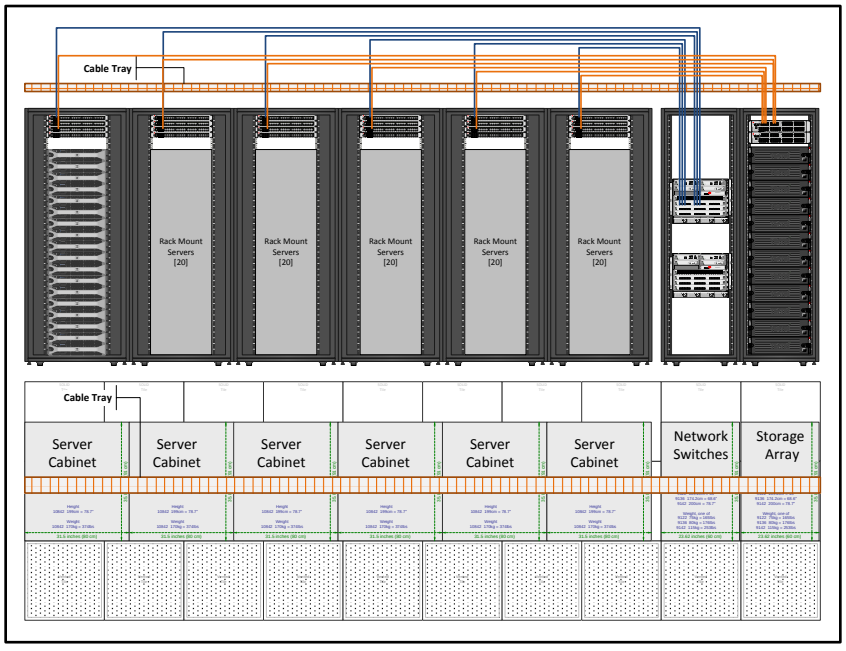


**Figure 13 - Virtual Infrastructure Rack Layout**

Figure 13 shows a possible rack elevation and floor layout for the virtual infrastructure design. This particular configuration would support over 2,000 VM's within each cabinet[9], overall number of virtual guests within six to eight cabinets, would be approximately 12 to 16 thousand. A storage cabinet with approximately 300 TB(s) of storage would be needed for every virtual cluster of servers, assuming each virtual guest, on average, required 80 to 100 GB of storage capacity. The configuration shown in the rack elevation would require at least two more cabinets of storage assuming these sizes are correct.

The major advantage to the approach outlined above is that each cabinet of servers can be attached to a particular cabinet of storage. As the server cabinets are filled with dedicated rack mount servers to support the virtual infrastructure, the storage capacity can be increased by adding storage modules to support the added servers. The correct balance between the servers and storage modules needs to be determined, though initially, the environment could be started with some less capacity in three cabinets (2 server cabinets and 1 storage cabinet) and then increased over time until an appropriate balanced ration of storage to virtual guests is realized.

### 3.2.3 Virtual (z/VM) Infrastructure

The virtual z/VM infrastructure maintains the same type of shared storage and shared network model as the prior two infrastructures detailed above. The primary difference is the scale or size of the server infrastructure (A mainframe versus a cluster of x86 servers) and the particular storage platform in use. For the z/VM virtual machines running on the mainframe several storage options exist, though IBM XIV (shown in Figure 14) is currently in use.

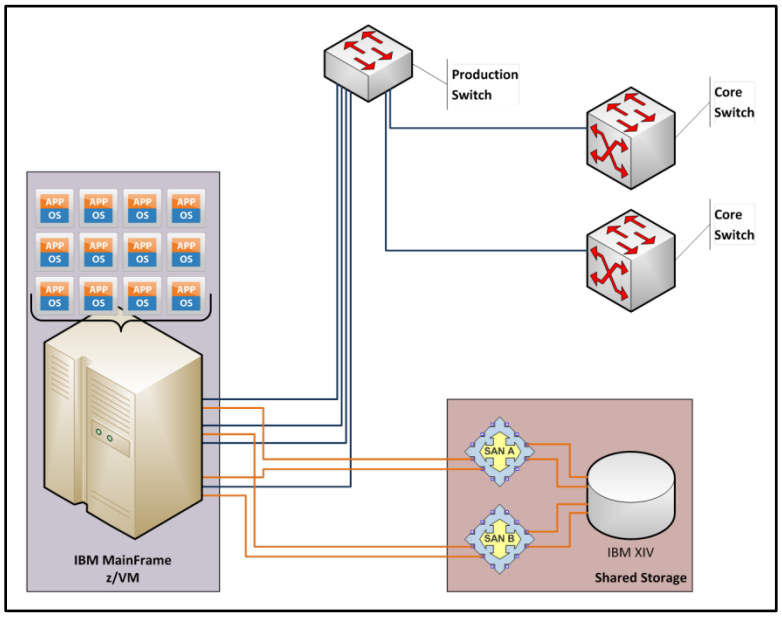


Figure 14 - Logical z/VM Virtual Infrastructure

The remaining connections to the mainframe z/VM environment and IBM XIV storage are similar as those for the physical x86 infrastructure, FC based storage network and Ethernet based network connections. The storage network is isolated from any other storage network, FC or Ethernet based, to

[9] This assumes a 4:1 oversubscription for vCPU and 4 GB memory assigned to each VM.

simplify the overall configuration – fewer components, less maintenance and less administration and any dependency with the other types of infrastructure has been eliminated.

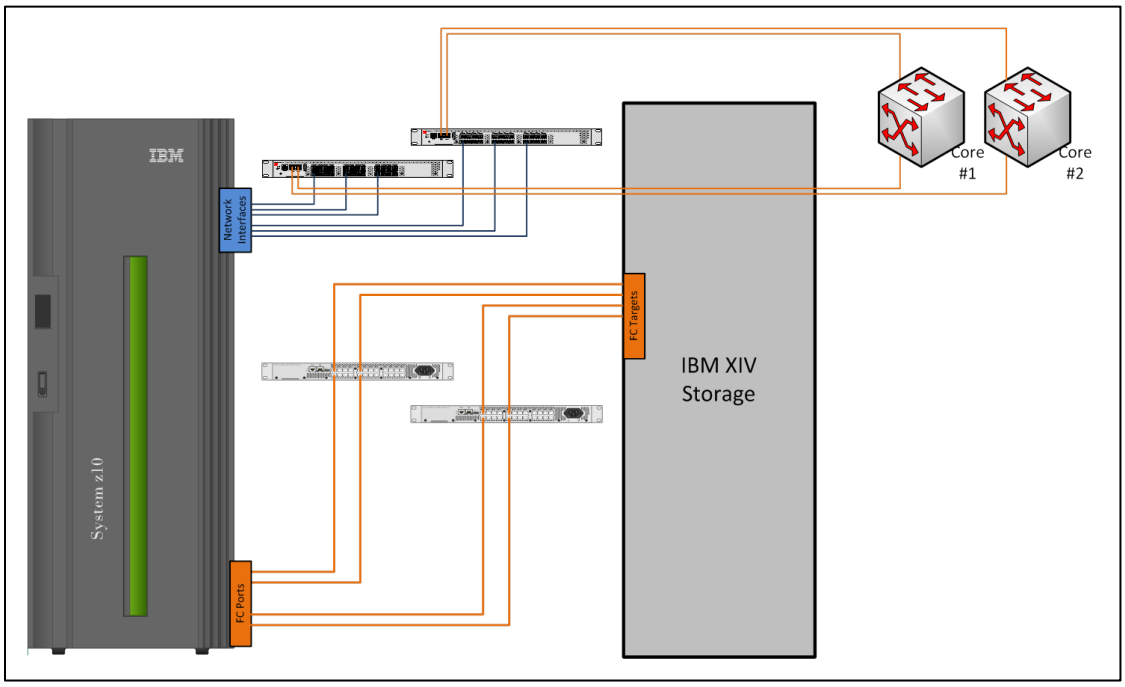


**Figure 15 – Cabinet Layout z/VM Virtual Infrastructure**

The cabinets shown in Figure 15, represents the entire environment with some room for growth, though current plans are to limit growth as much as possible. No cabinet elevation or floor layout is provided as the infrastructure is composed of only two cabinets and limited storage or IP network infrastructure.

## 3.3 Network Design

The network design is limited to the Ethernet switching infrastructure as it applies to all three types of infrastructure (virtual x86, virtual z/VM and physical) with the focus on virtual x86 infrastructures. The overall topology is shown below in Figure 16. This is considered to be a collapsed Access-Distribution-Core network topology since the distribution layer has been eliminated. It is also referred to as leaf-spine architecture with the core switches considered the spine and the access switches the leaves.

Although this is similar to current network infrastructures, the proposed network will be a partial meshed fabric[10], not shown are the inter-switch links (ISL) among the spine nodes or core switches. Six ISL 40 GbE links will connect each core switch to every other core switch with 2 40 GbE links to the each of the other three core switches. Each leaf node or access switch is connected to each core switch and each core switch is also connected to each router to insure that every core switch has an external connection for routing purposes. There are two leaf nodes or access switches at the top of each rack for redundancy.

[10] For a fully meshed fabric, additional links would exist between the leaf and spine nodes and links would exist between the leaves as well.

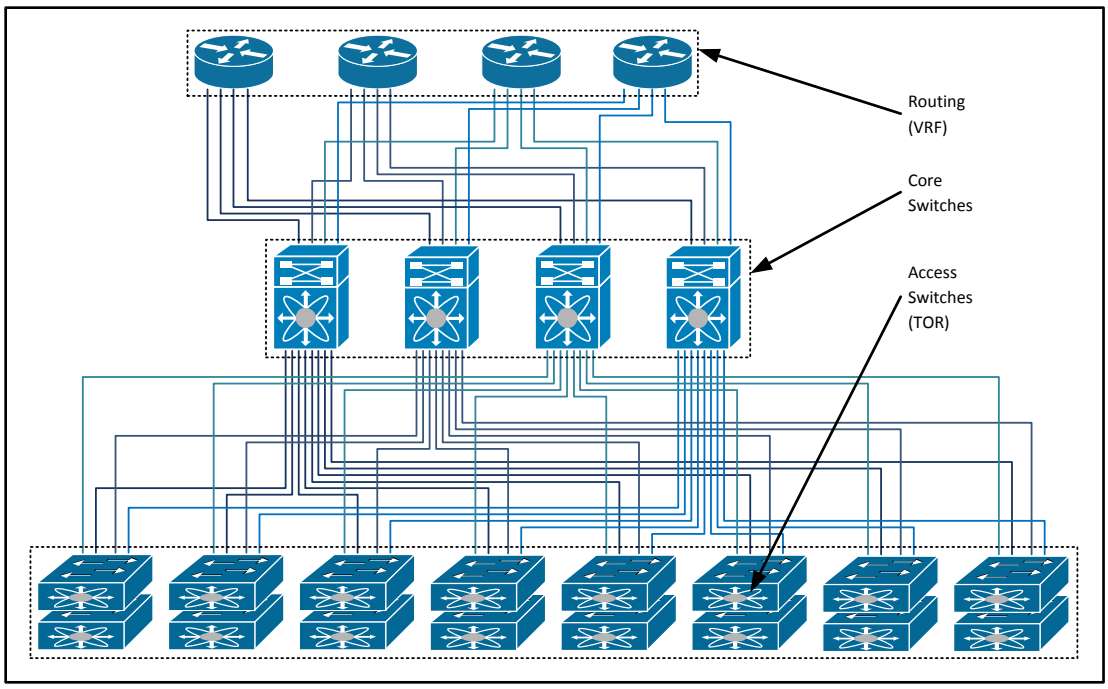


Figure 16 - General Network Topology

This is a diagram of the physical network and does not show the virtual networking components, which for the virtual x86 and Mainframe environments plays a role. It also does not show the x86-based virtual appliances that provide VPN functionality, traffic management, high availability, monitoring and reporting, administration and authentication, and firewall protection.

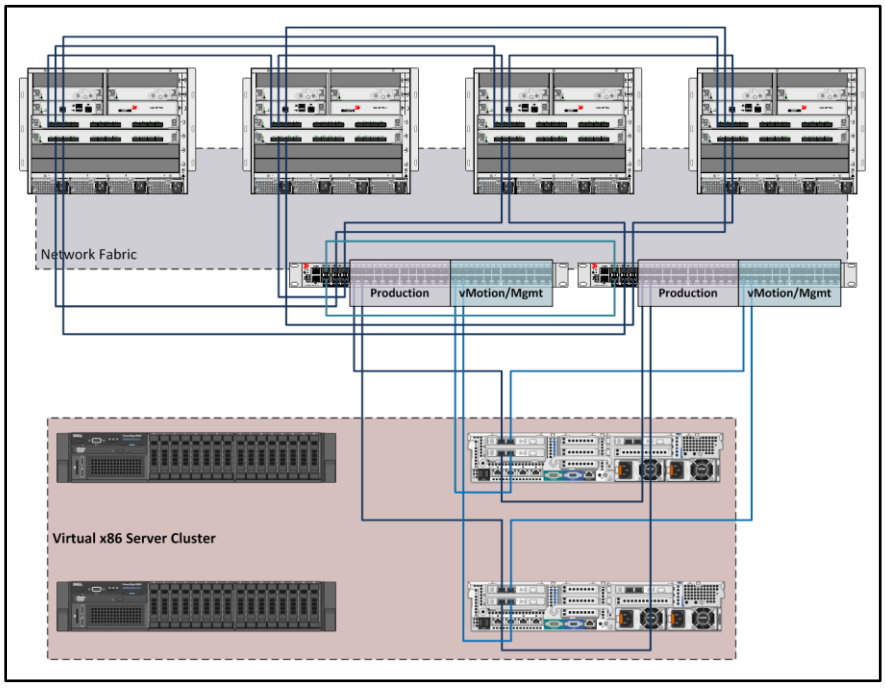


Figure 17 - Physical Topology for Virtual Infrastructure

The diagram above (See Figure 17) shows the connections and overall topology for (x86) servers supporting virtual infrastructure. Only a single connection from each core switch to every core switch is shown, though two should be implemented. Not shown are the inter-switch links[11] between the pair of access switches that would be located at the top of the rack. All of the connections from the server to the switches are at 10 GbE, with connections from the access switches to the core switches specified as 40 GbE. With the use of two 10 GbE switches each with 48 ports and assuming no more than 16 switches are installed in a cabinet, then each switch could support up to 3 connections from each server, though only two are shown. Another pair of connections could be implemented for either more bandwidth or additional networks such as a dedicated backup network or dedicated inter-VM network.

Although the two diagrams, Figure 16 and Figure 17, are different since the first shows the logical topology and the second shows the physical connections for only a subset of components, there is a clear relationship between the two. The diagram below (see Figure 18) depicts the virtual network connections from within the hyper-visor out to the physical network.

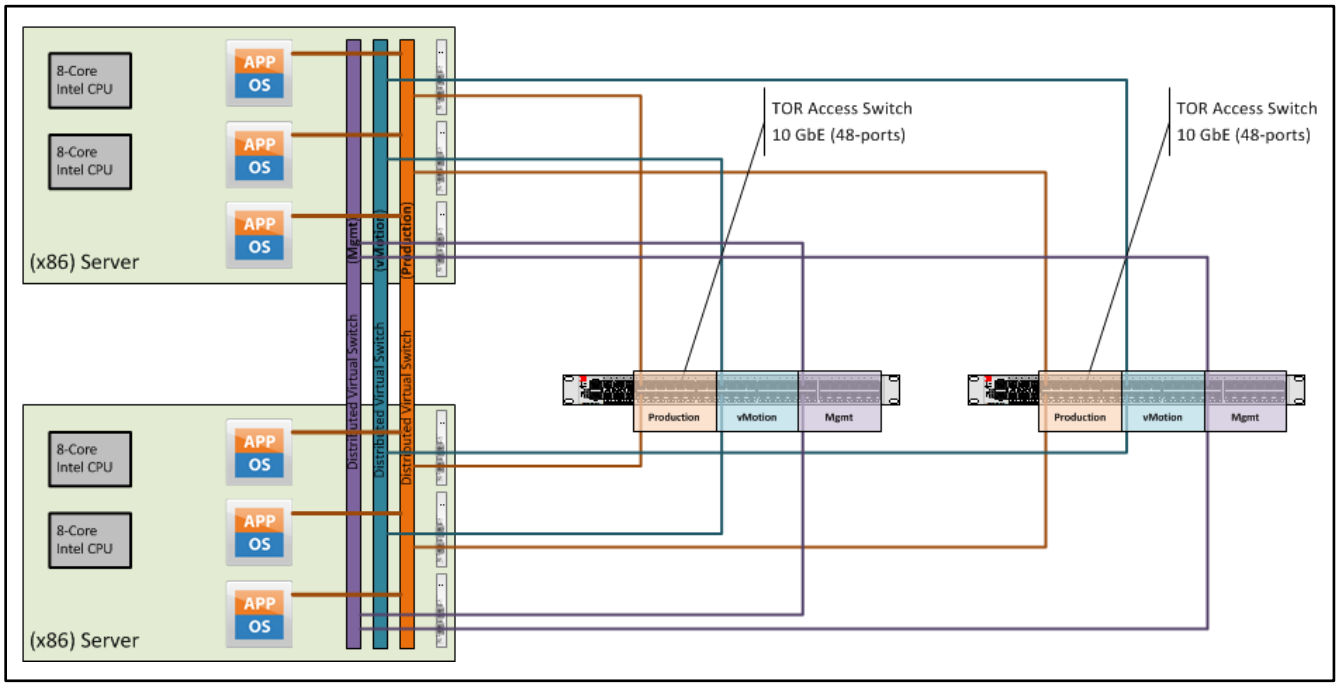


**Figure 18 - Virtual Network (Distributed Switches)**

Each distributed switch within the hyper-visor has two connections (one to each switch) from each virtual infrastructure host for redundancy. Each distributed switch is spread across all of the virtual hosts in a particular cluster within a cabinet. Each access switch at the top of every rack supports three major sub-networks, Production, vMotion, and Management. The Production network supports multiple VLAN(s) each created on the local access switches at the time the VM is provisioned. The vMotion network supports High-Availability (HA) and Dynamic Resource Scheduling (DRS) of the virtual guests within a particular cluster of virtual hosts. The management network provides additional capabilities such as a dedicated backup network link to offload traffic from the production network and potentially

[11] These may or not be implemented depending on the preference of the network team.

provide a separate network for network traffic between individual virtual guests. This would be provisioned with multiple VLAN(s) – one VLAN dedicated to each agency customer.

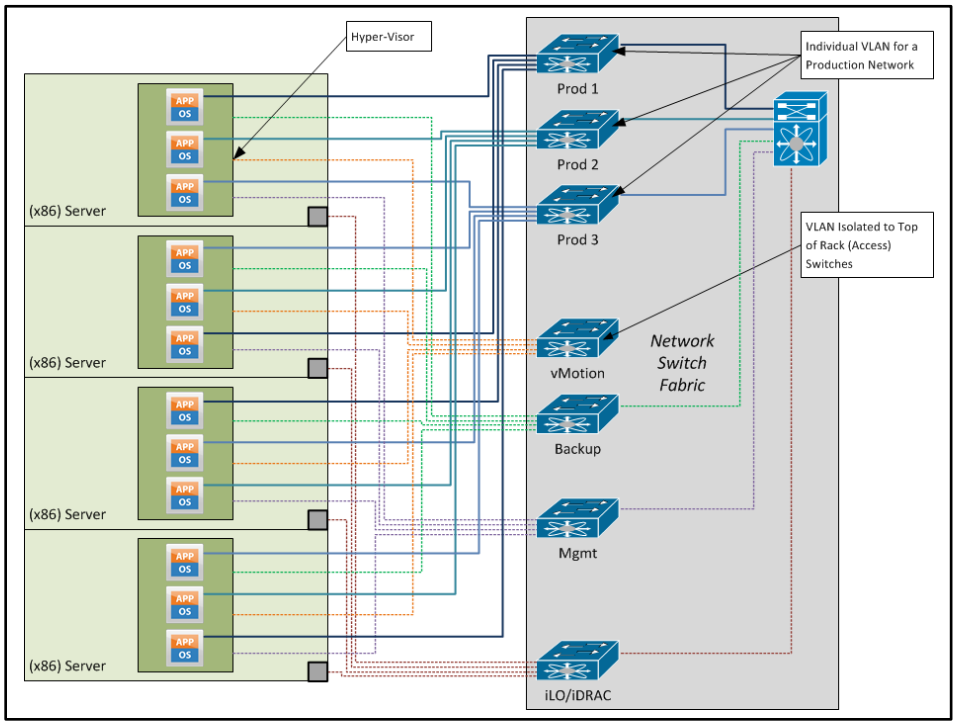


**Figure 19 - Logical VLAN Topology**

Although an agency customer might have multiple VM(s) running among several virtual hosts, each VM would be connected to a particular VLAN specified for that agency customer. In the diagram above (Figure 19), three VLAN(s) are shown with different virtual guests connected to each VLAN. Each virtual guest is connected to only one VLAN, though it could be connected to several VLAN(s) depending on the needs of the agency. All management, backup and production VLAN(s) are created within the entire network fabric, the exception, the vMotion VLAN, exists only on the access switches at the top of each cabinet.

There are a few differences between the logical VLAN topology and the virtual network diagrams. The primary differences are the number of VLAN(s). The virtual network diagram is intended to depict the distributed vSwitches. The production vSwitch would have multiple VLAN(s) created to support each customer environment – a completely isolated network environment. The logical VLAN topology is intended to depict the various types of connections and VLAN(s) in use within the network infrastructure, note that no distributed vSwitches are shown, but would exist between the virtual guests and logical switches within the network fabric.

## 3.4 Server Design

The server design is specifically for the design of the x86 physical servers supporting the virtual x86 infrastructure. This does also include the physical infrastructure supporting applications running on bare metal without a hyper-visor. These physical servers would be purchased based on pre-determined sizes, such as small, medium and large configurations that an agency customer could choose from when specifying their infrastructure. This design does not include the mainframe virtual infrastructure server

design as this has already been deployed and is in use. The mainframe also has far fewer choices when being configured for particular workloads and is well understood within the DAS/BEST infrastructure and therefore a design would be unnecessary at this time.

There are several vendor models capable of working within the virtual x86 infrastructure. This design will use a generic server and show how it would integrate into the network and storage infrastructure as shown in other sections. The general specifications for the generic server are that at least sixteen (16) servers can fit into each cabinet with sufficient room at the top of the rack to support multiple network switches for connection to the network switch fabric – this limits the height of each server to 2U. Each server must have at least four (4) CPU(s) with at least 8 cores for a total count of 32 cores for each server. The server must have at least three (3) PCIe 3.0 slots to support three (3) 10 GbE network interface cards, dual power supplies and two hot swappable SAS drives (preferably small 50 GB SSD drives). Additional network interfaces (1 GbE) should be available for use in managing each hyper-visor.

Each cabinet should be capable of supporting over 2,000 virtual guests. Therefore, given 16 to 20 servers in each cabinet, each server should be able to support 125 virtual guests or 4 guests[12] for every core. Using a standard VM size that requires 4 GB memory, each server would require 512 GB memory or eight (8) TB(s) of memory within the virtual infrastructure cluster in a single cabinet. Physically, this infrastructure would look similar to that shown in Figure 20.

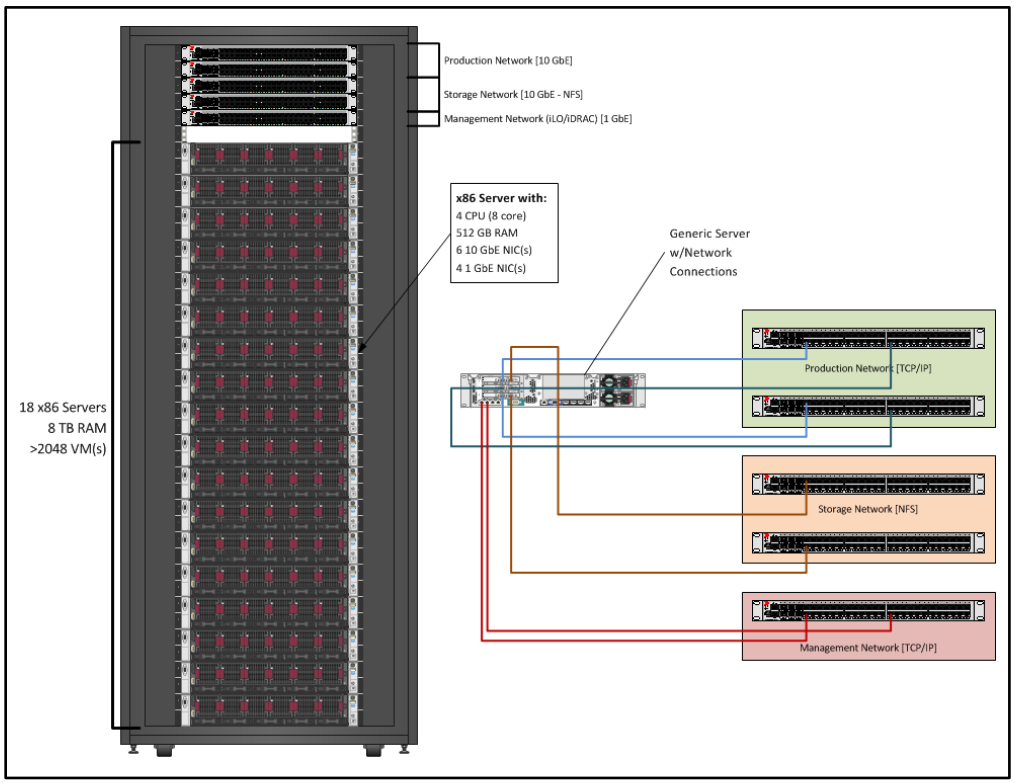


**Figure 20 - Server Design [Cabinet/Network]**

[12] We are assuming that each generic guest has 1 vCPU and 4 GB RAM allocated, with a 1:4 oversubscription ratio of guests to cores. We are also assuming that a guest requires a 50 GB root disk and another 100 GB disk for additional storage.

The switches at the top of the rack could be modified to support different needs of the overall virtual infrastructure depending on the specific needs of the DAS/BEST organization. For example, assuming enough ports existed on the access switches, the production and storage network could be collapsed onto a single pair of access switches. If additional management connections are required, a second management switch could be installed. An example of the generic server is depicted below in Figure 21, with the appropriate number of CPU(s), memory and network connections.

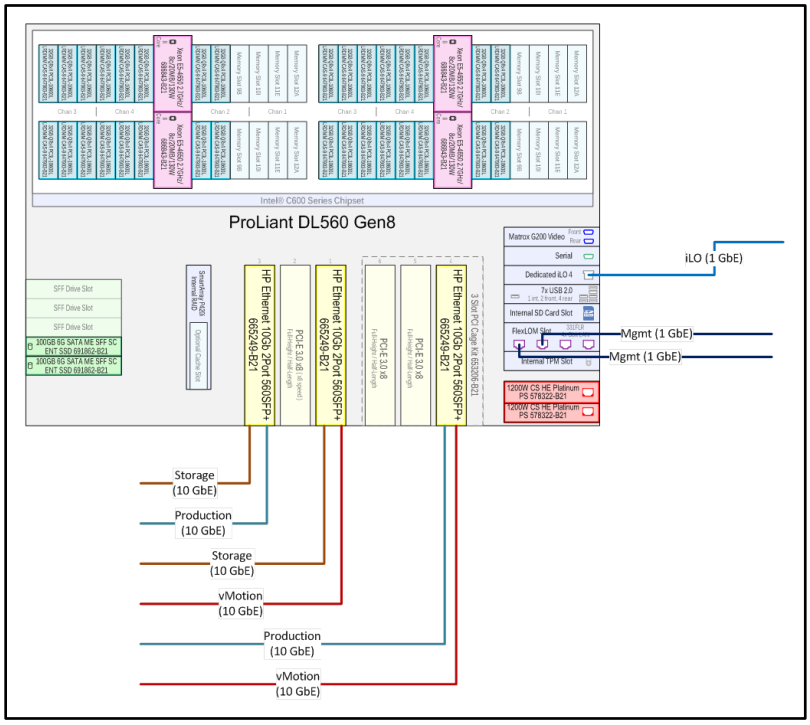


**Figure 21 - Generic Server Design**

Although the diagram shows the network connections dispersed across the three cards, it may be necessary to keep like connections on cards using the same PCIe version. It may also be possible to specify two 10 GbE ports for the onboard ports so that these could also operate at 10 GbE allowing either the elimination of one 10 GbE dual port card or the use of additional links for the production network. If the latter were chosen, the management network would be created on at least two of the virtual network distributed switches as described in the prior section.

## 3.5 Storage Design

The storage design has three major areas. These are the storage design for the virtual (x86), virtual (z/VM) and physical (x86) infrastructure. Based on the initial options for infrastructure models, the hybrid/isolated model will be used. This means that storage will be isolated by infrastructure type – a separate storage infrastructure for virtual (x86), virtual (z/VM) and physical infrastructure.

Each infrastructure will also use a different storage and network protocol. The virtual (x86) infrastructure will use a Ethernet network and the NFS protocol (a file based protocol), while the other two types will use a Fibre-Channel based storage network and shared storage based on a block protocol (SCSI) instead of a file protocol (NFS). Each type of infrastructure will be isolated from the other infrastructures, such that the storage cannot be shared among the various infrastructure types. The

virtual (x86) infrastructure will not be able to use the storage dedicated to the virtual (z/VM) infrastructure. Each infrastructure will have a specific amount of provisioned capacity with future growth expected. Table 1Table 13 provides an approximate breakdown of these environments, the storage in use, the initial capacity and potential future growth.

| Infrastructure | Storage Protocol | Network | Initial Capacity | Growth (2 Years) |
|---|---|---|---|---|
| Virtual (x86) | NFS | Ethernet | ~170 TB(s) | 170 TB(s) |
| Virtual (z/VM) | SCSI | FC | N/A | N/A |
| Physical (x86) | SCSI | FC | ~200 TB(s) | 50 TB(s) |

Table 13 - Capacity Estimates

### 3.5.1 Physical (x86) Infrastructure

The storage design for the physical (x86) infrastructure will not change appreciably from its current implementation. Currently the physical storage infrastructure (SAN) for physical servers (operating environments running on bare metal) is based on FC based storage switches and FC attached storage arrays. The major change to this infrastructure is the use of shared storage using fully automated storage tiers within the array to provide the best performance possible to the applications sharing the storage. In addition, thin provisioning will be used to reduce initial actual capacity assigned to hosts – storage would grow as needed for an individual application.

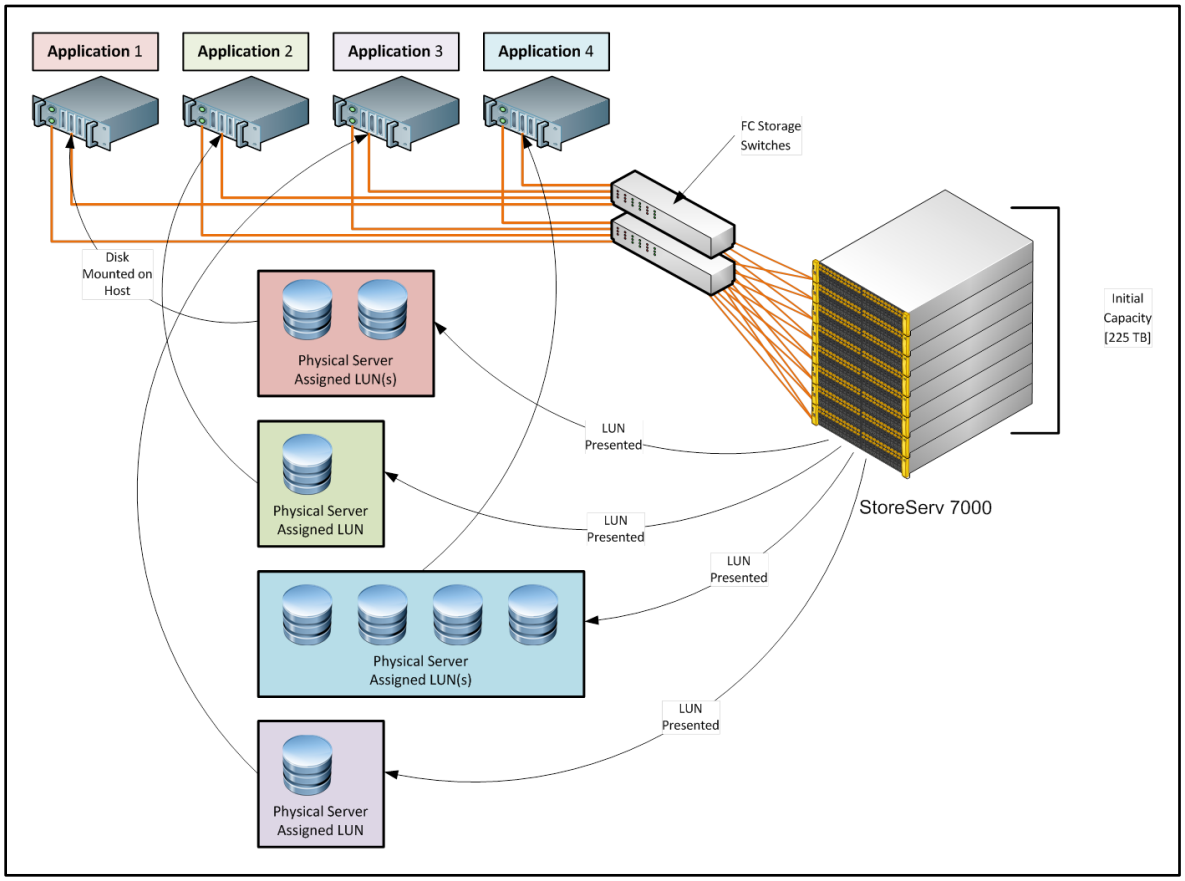


Figure 22 - Physical (x86) Storage Infrastructure

The storage would remain FC based using the same infrastructure as is currently in place. The primary difference is the change in storage array model and vendor with some substantial consolidation provided by newer technology.

Increasing the capacity for the shared storage in the physical (x86) infrastructure only requires purchasing additional capacity until the limit of an array has been reached and then the purchase of a second array to meet ongoing storage needs. Alternatively, two arrays could be purchased with sufficient capacity to meet current needs, enough spare capacity to meet growth needs for two years (the standard state budget cycle) and enough spare room in empty slots for future growth. The spare room for growth can also be provided with setting spare room within the cabinet as well, since most if not all storage arrays allow for the expansion of growth by adding storage trays.

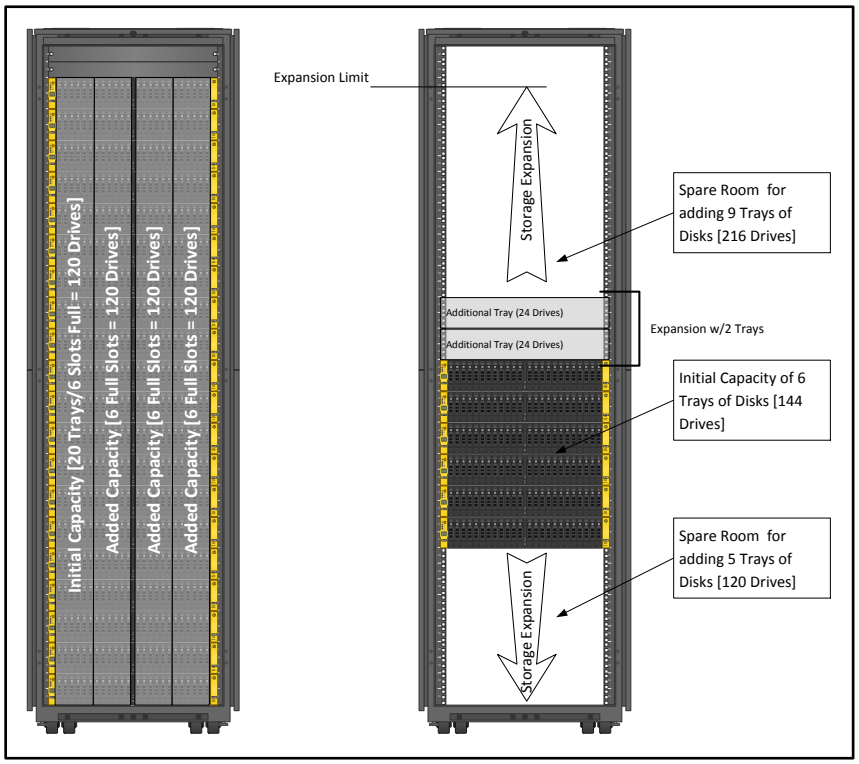


**Figure 23 - Storage Expansion Options**

Figure 23 above depicts the two options for expansion with two example cabinets. The left hand side cabinet is full of disk trays, drives are installed in a single column across all of the trays within a cabinet. Adding capacity is achieved by adding more disks in one or more columns in all of the trays of the cabinet. The right hand side cabinet shows the cabinet with only six trays installed and the initial controllers. These trays can be full of drives or partially full. Adding additional trays with drives provides added capacity until the cabinet is full. Both approaches are valid and both work with almost all mid-range storage arrays including those from IBM, HP and EMC.

Additional design for this infrastructure would be based on host requirements for storage capacity and performance. Storage array design is based on configurations that provide the most capacity and the best performance. For any given array, the physical hardware (composed of controllers, trays and disks)

is configured to spread the data blocks used by a particular application across as many components (usually disk spindles) within the array as possible. Using fully automated storage tiers (FAST[13]) within an array, this spread of blocks within the array is accomplished automatically, always providing the best performance to all applications.

### 3.5.2 Virtual (x86) Infrastructure

The storage design for the virtual (x86) infrastructure is intended to provide a simple, highly redundant, high performance storage infrastructure to support the virtual guests so that any (x86) whether Windows or Linux workload, can be placed in the virtual infrastructure. The logical topology of the storage infrastructure is shown in Figure 24 below. Each modular storage unit presents a single 33.5 TB NFS mount over 10 GbE interfaces[14] connected to dual 10 GbE switches to the virtual infrastructure.

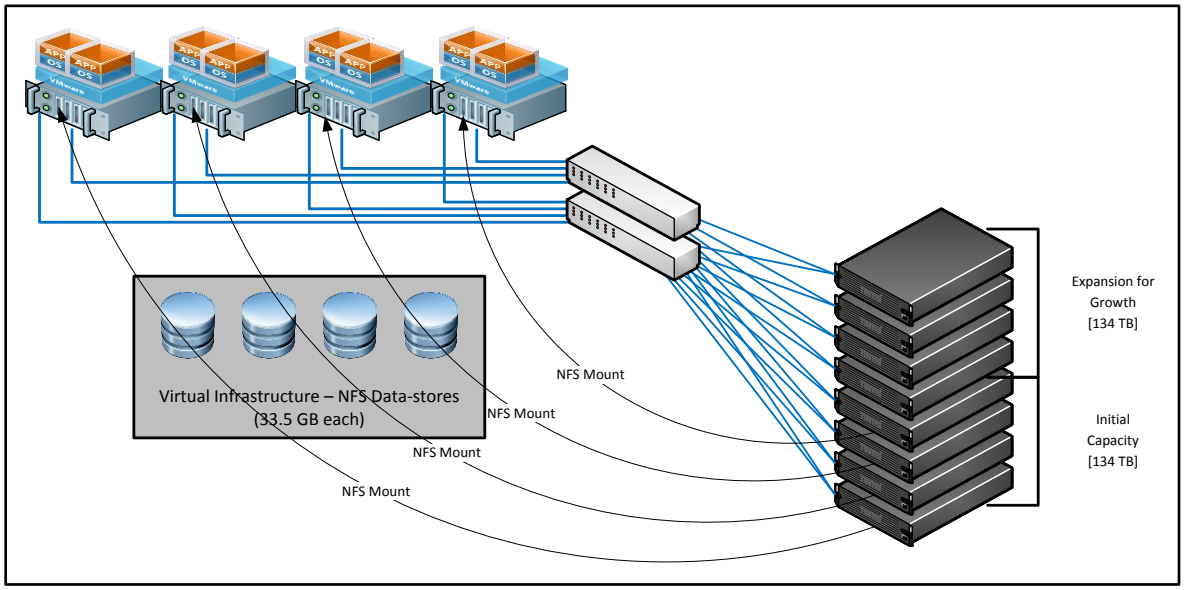


Figure 24 - Virtual (x86) Storage Infrastructure

The virtual infrastructure hyper-visor is VMware vSphere v5.5 ESXi. The maximum allowed number of NFS mounts for a particular host is 256[15] or, in this case, over 8 Petabytes of storage can be allocated to the hosts before the number of mounts are exceeded. The other limitation is the number of storage units that can be managed as one. Currently, for this storage model, that limit is roughly 30 units or slightly over 1 Petabyte of storage capacity for the virtual infrastructure in 3 cabinets.

The physical design for the storage includes the storage network and host connections required to support access to the modular storage units using the product from Tintri – the T640 storage array with 33.5 TB of useable storage accessed by the VMware ESXi using the NFS protocol. The storage network uses TCP/IP over Ethernet with quality of service (QoS) set to insure reliable traffic for all NFS packets. Figure 25 shows the overall storage and storage network for the virtual infrastructure. At the core of the

[13] Different vendors call this by different acronyms. Almost all vendors provide a feature that offers this functionality, typically based on license fees.
[14] Either SFP+ or RJ-45.
[15] The default is number of NFS mounts allows is eight, 256 is the limit if the advanced NFS mount setting is changed.

design is a two-tier switch infrastructure. The lower pair of switches aggregates the storage nodes and makes all nodes accessible to all hosts in the network.

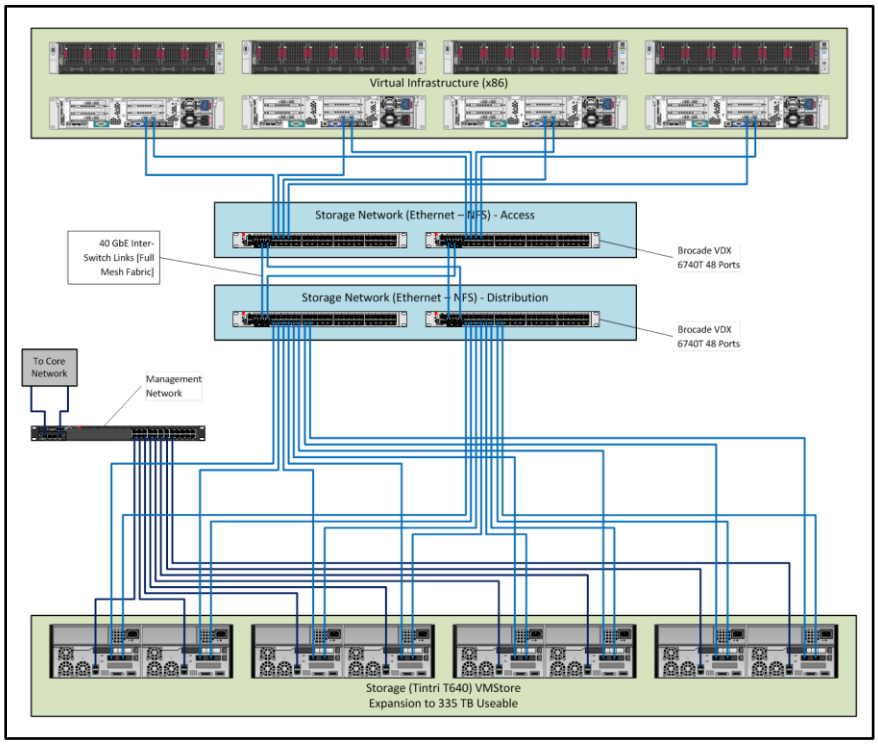


Figure 25 - Storage Network for Virtual (x86) Infrastructure

Both pairs of switches are interconnected in a full mesh fabric with all links active and balanced such that the same percentage of traffic travels over each link – no link is over utilized compared to the others. This particular approach assumes that a single cabinet of storage (roughly 335 TB Useable) will support at most two cabinets of virtual host servers, which means that for 4,000 virtual guests (2,000 for each cabinet); each guest will require or use approximately 85 GB of storage. If the required storage is more or less than 85 GB/VM, then some adjustment needs to be made to either the number of storage nodes or the number of virtual hosts that are allowed access or to the type of switches used for the storage tier. Two approaches are available to either size the virtual infrastructure or to scale the storage network as needed and are represented in Table 14 below.

| Approach | Switches Used | Number of Hosts | Number of Storage Nodes | Network Topology |
|---|---|---|---|---|
| Scale Servers/Storage | Brocade VDX 6740T 48-Port 10 GbE (Quantity 4) | 16 Virtual Hosts (ESXi)<br>1 Cabinet | 9 Units<br>1 Cabinet | Use full mesh between all 4 switches |
| Scale Network | Brocade VDX 8770-4 w/48-Port 10 GbE Modules and 12-Port 40 GbE Modules | 192 Virtual Hosts (ESXi)<br>12 Cabinets | 108 Units<br>12 Cabinets | Use partial mesh to core switches.<br>2 core switches, with access switches for storage and virtual hosts |

Table 14 - Storage Network Topology Options

Although the virtual hosts running VMware ESXi support only 8 NFS mounts as a default, this can easily be expanded by modifying one of the advanced settings that control this number. However, if the initial approach taken is to scale the servers and storage together to the maximum number of components supported by the TOR switches, then a choice can be made to reduce the number of nodes in a single cabinet to eight and attach these storage nodes to a single virtual cluster or set of 16 virtual hosts. This approach can be taken in either case; it is simply a matter of choosing the appropriate network topology. Ultimately, the choice to scale the network means using a pair of core switches to support connections from multiple cabinets each with a pair of access switches to allow communication among many storage nodes and virtual hosts.

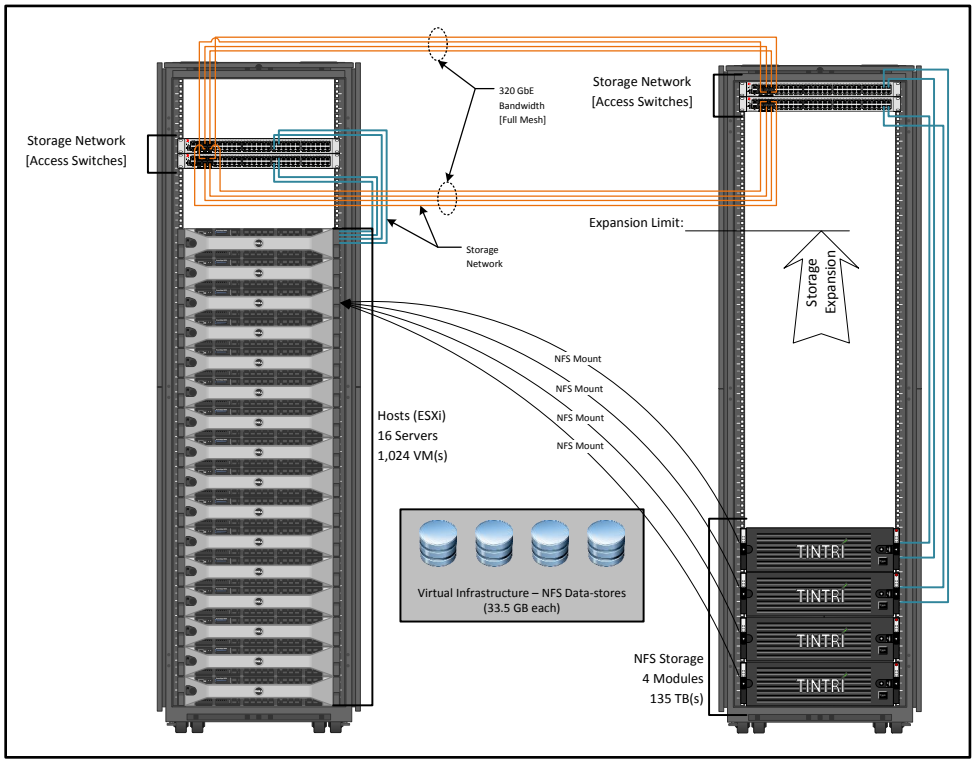


**Figure 26 - Virtual Infrastructure [Scaling Storage]**

Figure 26 provides an example of the first approach – scaling the storage as needed. The storage network is contained within these two cabinets and is a full mesh fabric with no oversubscribed links. A further reduction in infrastructure could be made by using on a single pair of switches and cabling the storage nodes directly to the switches supporting the virtual hosts. Assuming 16 hosts and 8 storage nodes, only 32 ports (each storage node requires two connections on switch) are required on each of two switches for all network interfaces to connect. Each NFS storage unit is presented to all 16 virtual hosts until either the default maximum of eight (8) NFS mounts is reached or sufficient storage is assigned to reach 75% utilization within the virtual hosts. If more storage is required, the NFS mount advanced setting is modified and further storage units can be added to the infrastructure.

## 3.6 Backup Design

The backup design is intended to provide a complete end-to-end solution for backing up all of the environments that currently exist at the DAS/BEST site. This requires that all three types of infrastructures have a backup solution in-place and includes virtual (x86) infrastructure, virtual (z/VM)

infrastructure, and physical (x86) infrastructure. The two primary operating system types in use in all three infrastructures are Windows (various versions) and Linux (various versions and distributions).

The design must scale to backup all of these infrastructures and provide a quick means to restore data as quickly as possible. There are two components to this solution. The first is the underlying hardware to support the backup environment and the second is the actual backup software that will be used to accomplish the backup and restore of data. Figure 27 depicts a logical view of the overall backup environment and the integration of the backup environment into the three primary types of infrastructure.

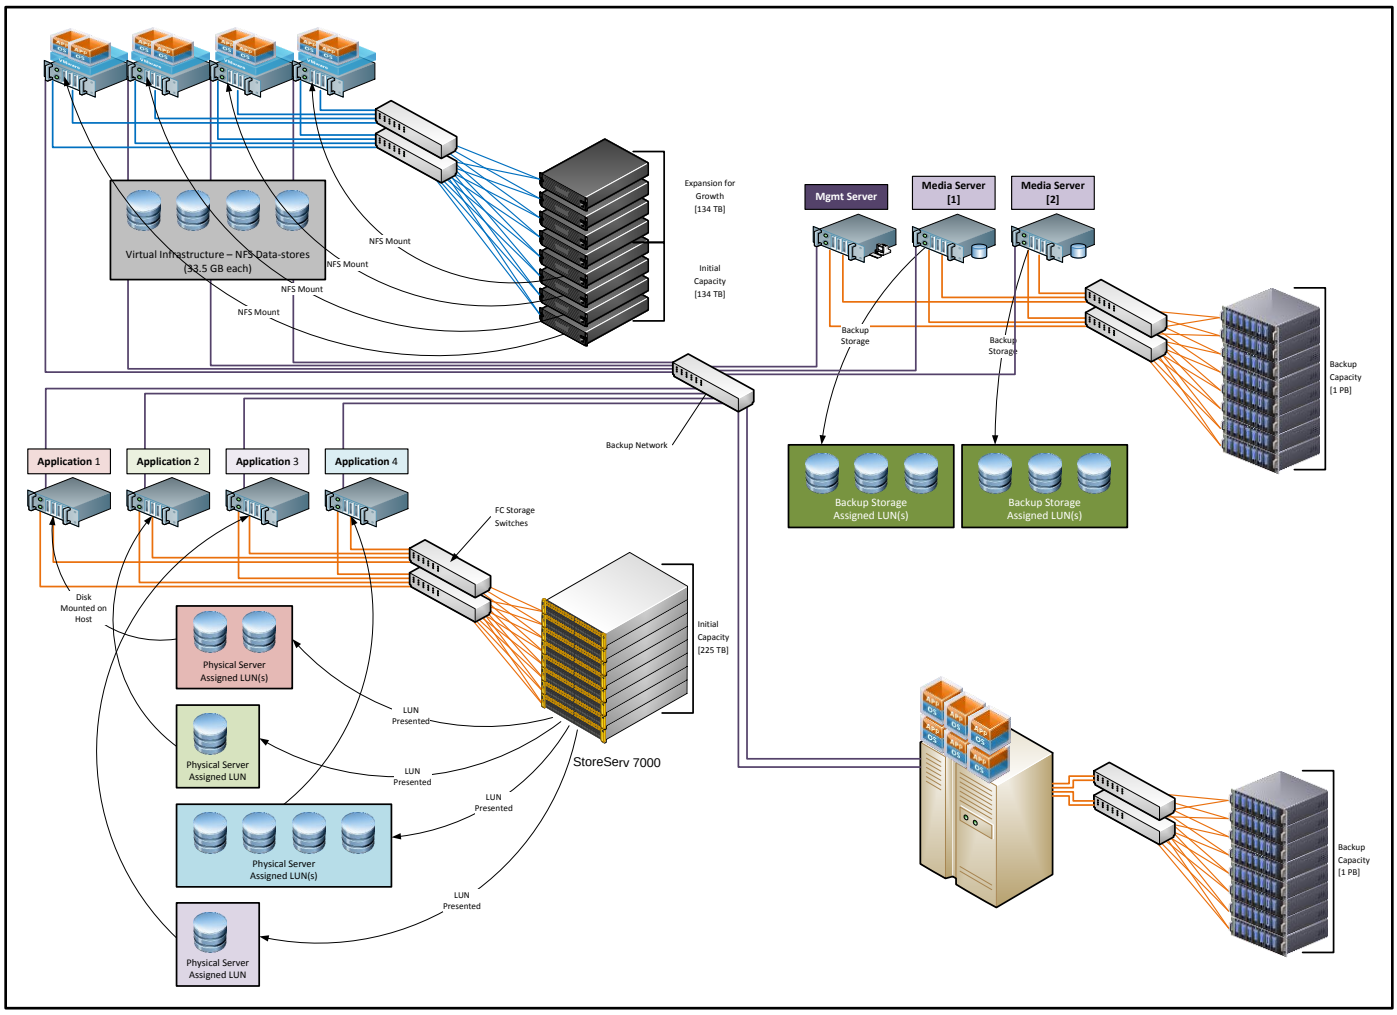


**Figure 27 - Backup Design Topology**

There are specific requirements for the backup solution and therefore aspects that the backup design must provide. The more important requirements are:

1. Snapshot integration with physical storage infrastructure to support snapshot based backups if needed
2. Virtual Hyper-Visor integration for virtual machine backups
3. Workflow automation for backup environment
4. Self-service access for agency customers requesting the ability to backup their own environment
5. Backup data de-duplication at source

6. Data replication to a remote/secondary site using de-duplication

Other requirements include the ability to manage the environment from a single point of control, scale the backup solution as needed over time using dedicated backup media servers, support of both physical and virtual environments and multiple versions of Microsoft Windows and various distributions of Linux. The final backup design includes the use of a single backup software product and a single, scalable, backup storage target to meet these requirements.

The backup hardware infrastructure is composed of several parts. The first and most important is the storage repository that will be used to store the backup data; to meet DAS/BEST backup requirements disk drives will be used to store this data. At this time, it has not been determined how much data must be backed up[16] though the standard retention periods and full/incremental schedules certainly apply.

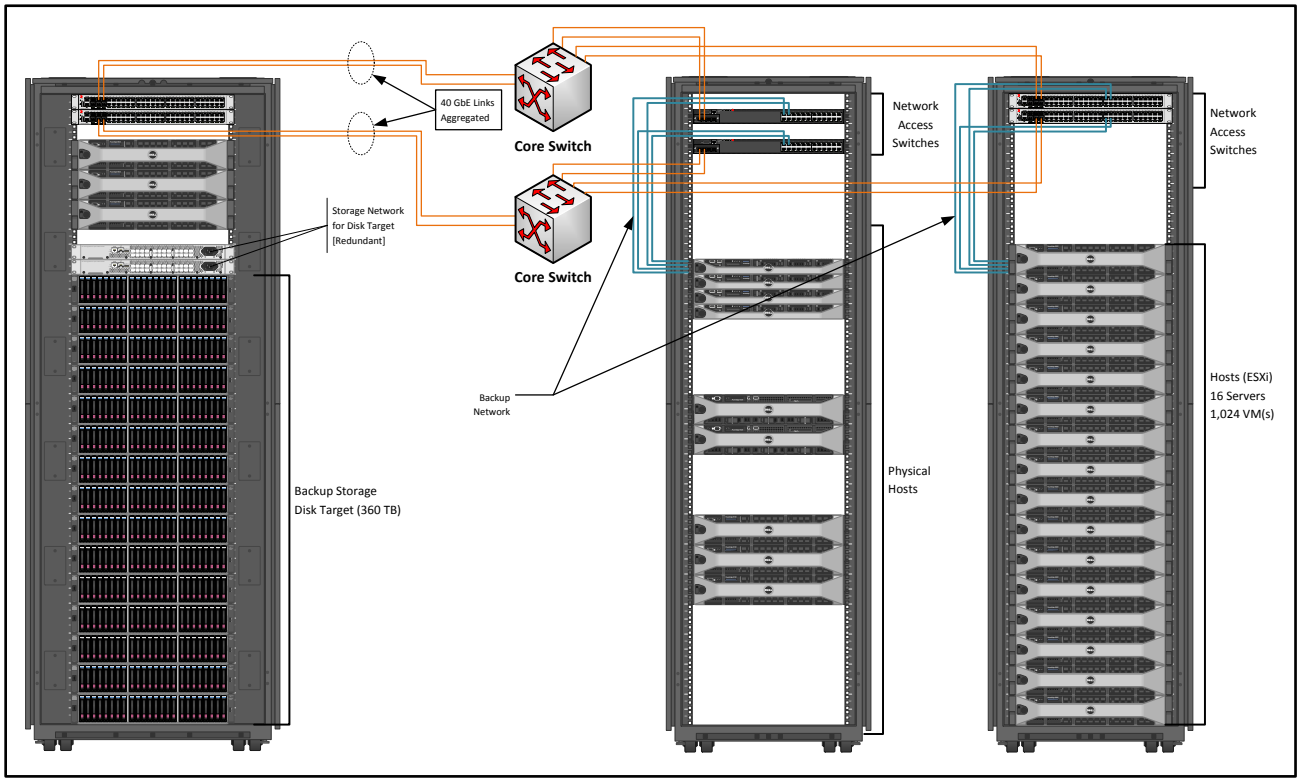


**Figure 28 - Backup Hardware Infrastructure**

Shown in the diagram above (Figure 28) is the overall backup infrastructure – servers shown to provide detail of integration with network. The storage target shown is a collection of 15 HP MSA P2000 G3 Storage arrays with a combined raw capacity of 360 TB's. Utilizing RAID-5, this would become, roughly, 280 TB of useable capacity. Each array would present a single LUN to the backup media servers to store backup data. The local network used would be either a FC or an iSCSI based storage network. Other storage disk targets include the Exagrid and Data Domain storage arrays. These arrays use the NFS protocol as the primary means of communication. These arrays would be on an isolated network accessible only by the backup servers.

[16] Current storage used is roughly at 350 TB's.

The Exagrid array has a 21 TB model that occupies four rack units (EX21000E). The overall grid can scale out to 10 nodes or 210 TB(s) accessible as a single unit by the backup software. Data Domain uses a scale up approach to increase capacity, and also provides access using the NFS protocol as well as FC based Virtual Tape Library (VTL).

The backup infrastructure is designed to use a dedicated network for all virtual hyper-visor based infrastructure and guests, with either a dedicated or a shared network for all physical servers. The backup network used for the virtual infrastructure would not be accessible by the guests, only the hyper-visor. These two infrastructures should be isolated on different VLAN(s) so that no communication is possible from the backup network interface used by the physical server to access the virtual infrastructure. If the physical servers will share the production interface or other interface for backup traffic, this will need to be routed to a dedicated VLAN for only physical server backup traffic. These VLAN(s) would be connected to the backup servers on different links, one for the virtual infrastructure and one for the physical infrastructure.

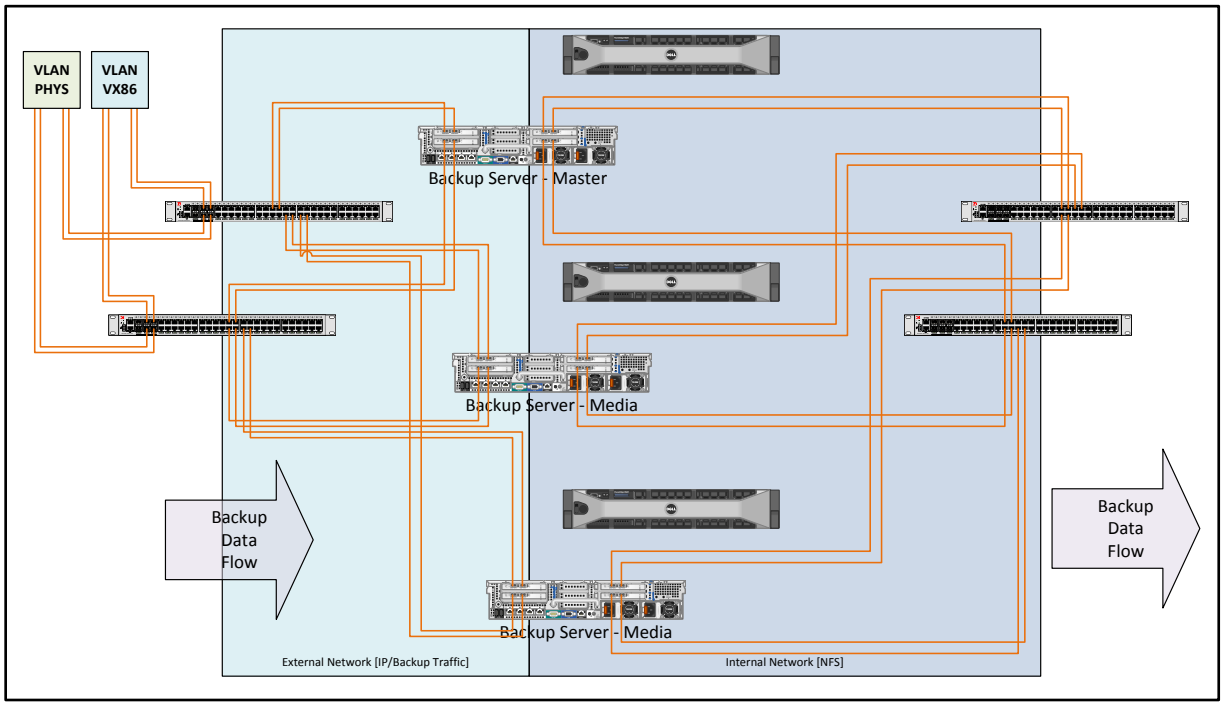


Figure 29 - Server Network Connections

The backup infrastructure would scale by adding additional nodes or media servers to the environment to backup data from the various virtual and physical servers to the storage disk target. Ideally, the backup infrastructure must be scaled to support the amount of data within a proscribed backup window. For example, if 60 TB of incremental data must be stored every night, and the window is six hours, then the rate is 10 TB/hour. If 10 GbE Network links are used within the servers and each link supports data rates of 2,000 Mbps, then each 10 GbE network interface would have a throughput of 6.7 TB/hour.

To provide a throughput rate of 10 TB/hour, at least two interfaces must be used on the inbound side and two interfaces (either 8 Gbps FC or 10 GbE) used on the outbound side. In addition, the storage repository must support receiving this data at the same rate. Assuming the use of Exagrid storage (model EX21000E) the maximum throughput is 4.32 TB/hour, so two units would satisfy the throughput

rate, though additional storage would certainly be required to support the amount of data to protect. Figure 30 shows the throughput for a single backup media server, assuming four incoming links each at 10 GbE and four outgoing links each at 10 GbE. This would be approximately the same throughput diagram assuming the use of 8 Gbps FC links to the storage, though there might be some slight differences in the throughput to the storage.

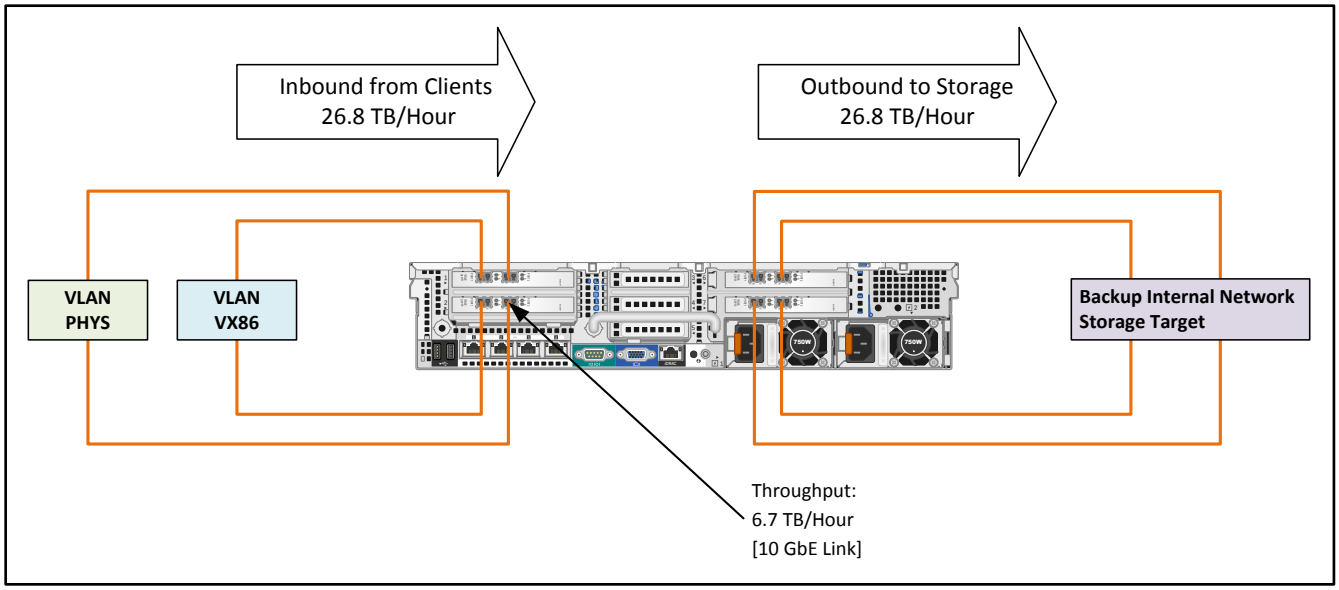


**Figure 30 - Data Throughput [Backup Server]**

## 3.7 Disaster Recovery Design

The disaster recovery design is intended to take advantage of an active/active data center model (see Section 3.1.1 Active/Active Data Centers) and provides a mechanism for agency customers to failover/recover their applications at a secondary site. Since both of these data-centers are intended to be in use, so applications would run at one site and failover to the alternate sites.

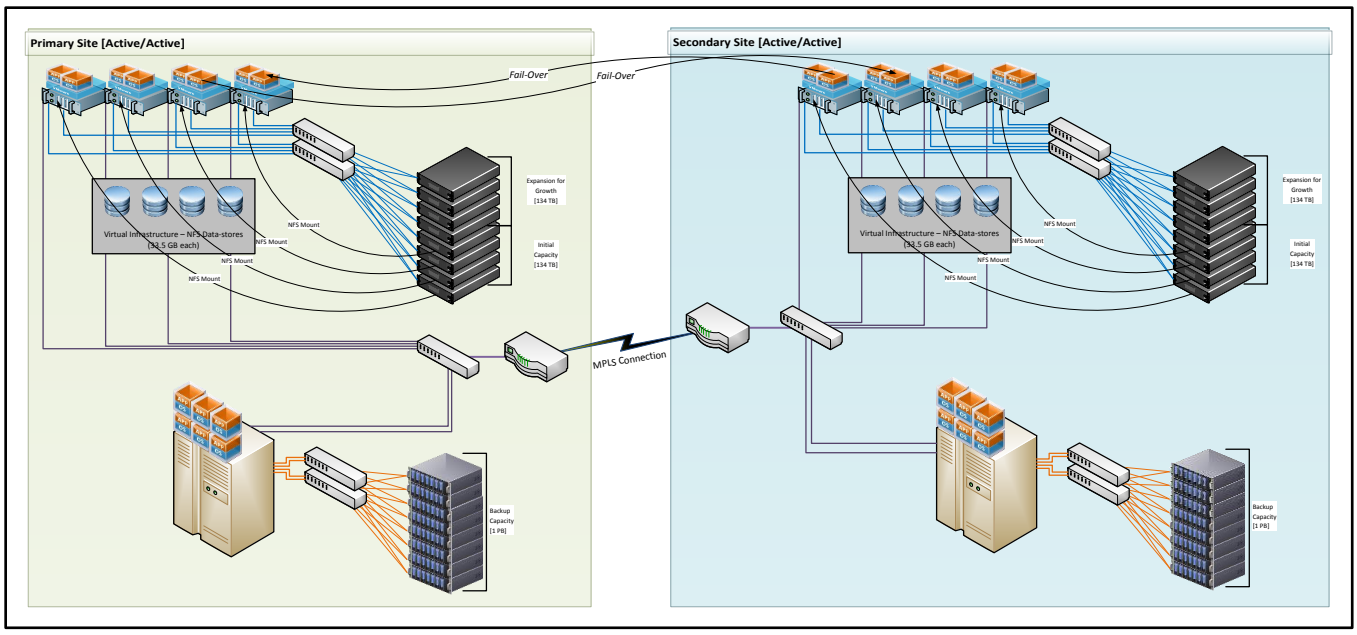


**Figure 31 - DR Topology [Virtual Infrastructure]**

Both sites would host active production applications with the alternate site acting as the secondary site (failover) for that particular application. Figure 31 above shows the overall logical topology for the two datacenters assuming virtual infrastructure is in use.

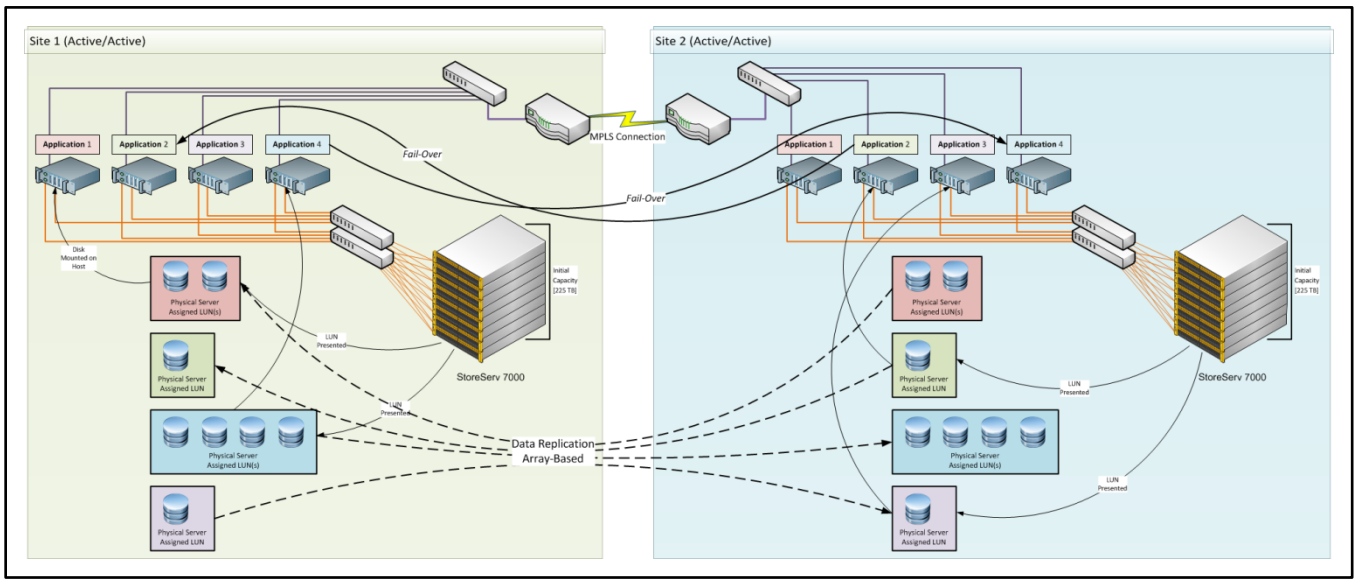


**Figure 32 – DR Topology [Physical Infrastructure]**

Figure 32 above shows the disaster recovery topology assuming physical infrastructure. There are several important differences between the two approaches. In the virtual infrastructure, data replication between the storage at either site is not required; replication happens at the hyper-visor layer, below the physical application. For physical infrastructure, this is not possible since no hyper-visor exists. Therefore, the underlying block data within the array must be replicated to the second site for future recovery of the application. This is seen in Figure 32, where each set of LUN(s) assigned to an application

is replicated to the opposite site. Make note that some applications are shown active at Site 1, some at Site 2, and that data is replicated by application and the associated storage to the opposite site.

In the case of physical infrastructure, no other tools exist to assist with the actual recovery of the application at the remote site. In order to complete the recovery, a specialized script needs to be written to bring the application to a running state.

In the case of virtual infrastructure, the design for active/active data centers supporting disaster recovery services simply requires an application that supports replicating virtual guests from one site to another. Currently the simplest, lowest cost approach utilizes Zerto running within the virtual environment to replicate each virtual guest and provide a simple set of automatically executed steps to bring the application online. For physical infrastructure, automation of the online recovery of each application must be custom built to execute the necessary steps to recover a particular application.

## 3.8 Security Design

The security controls that DAS/BEST manages for its' customers can be broken down into the list below:

- Physical Data Center access security
- Change Control
- Security Configuration Management
- Firewall
- Directory Services
- Multi-factor Authentication
- Intrusion Prevention System (IPS)
- Identity and Access Management
- Email Gateway
- Internet Content Filtering
- System incident and Event Manager (SIEM)
- Network device security
- VPN Remote Access
- Anti-virus, Anti-malware, Hard Drive encryption
- Application Security

Deployment options will be assumed from the recommended choices from the Infrastructure Options document delivered to DAS/BEST earlier in the SHI project. Strategy for deploying these options will be based on anticipated internal resource availability and technical dependencies between solutions.

### 3.8.1 Physical Data Center access security

Physical access security deployment should be carried out during construction on the facility while access to areas for cabling, electric and equipment mounting can be performed efficiently. Biometric multi-factor authentication readers need to be located where video surveillance can be utilized to also verify someone's identity. Installation and configuration for physical access control can and should be done as the first measure of strategy prior to installing or moving any infrastructure systems into the new Data Center. The only dependencies related to Physical Data Center Access Controls involve timing around construction schedules for installing the hardware and cabling. User accounts can be created locally and matched to biometric registration until IBM Security Identity Manager can link account registration through Active Directory. AD account linkage for biometric access can be made through IBM Security Access Manager when available.

### 3.8.2 Change Control

Change Control process enhancements can be made immediately without any dependencies on any other solution initiatives. Priority effort should be made to this security control since it will greatly reduce adverse production impacts in follow-on solution improvement initiatives. Inviting senior architects and directors to develop and improve Change Control standards and procedures would be a great place to start with designing an effective process. A process needs to be developed and approved immediately for Change Control around security control initiatives along with associated Data Center infrastructure initiatives.

### 3.8.3 Security Configuration Management

Tripwire can be deployed immediately and may assist with Change Control process streamlining. Consulting services engagement is highly recommended for DAS/BEST for Tripwire as Security configuration management along with Change Control process improvements will gain a large operational efficiency advantage.

### 3.8.4 Firewall

Firewall changes and additional VPN capability will require significant resources from the firewall security team. Engaging a Check Point consultant will provide the benefit of a streamlined rule base and prepare the overall firewall environment for VPN functionality. Cleanup and review of the policy and rules base should be completed prior to any new security blades such as VPN and additional agency firewall management responsibility. Streamlining rules base and additional changes to the firewall architecture, application support, and software blade functionality would likely require Change Control review and approval. Completion of primary Change Control process review and endorsement by senior management should be done prior to any major firewall structure or functional changes. Completion of the Network Address Translation project to remove internal public IP networks also requires Change Control and resources from the Firewall and Network teams. This NAT project, although minor, will demand planning and coordination which places additional emphasis on team coordination and Change Control.

### 3.8.5 Directory Services

DAS/BEST operates and maintains Microsoft Active Directory services for central authentication, which is leverages by systems and applications to permit access for authorized users. There is no change in operation of Active Directory moving forward but as usual, this service will become even more relied on as additional security controls such as IBM Security Identity Management and Biometric physical access integrate with Active Directory LDAP. Inevitably, more applications and services will look to Active Directory for user identification, which places a large reliance on maintaining the availability and redundancy of Domain controllers and the networks supporting them. There is no dependency on other processes except Change Control should additional domain, forests or domain controllers need to be added or decommissioned.

### 3.8.6 Multi-factor Authentication

Migration from RSA hard tokens to Symantec VIP soft tokens for multi-factor authentication will require Change Control to add a cloud based administration tools from Symantec and decommissioning RSA authentication server when all hard tokens have been replaced by soft tokens. Systems and applications that current authorize users through RSA tokens will also need to be reconfigured to authorize users through their new Symantec soft tokens. No other project dependencies exist but this move to soft

tokens will need to be coordinated with the RSA administration team and all applications and systems that user RSA tokens for authorized access.

Integrating Biometric multi-factor authentication with Physical Access to Data Center facilities will need to take place during the design and installation of the access systems at the new building. No other dependencies exist directly with Biometric authentication for physical data center access. IBM SIM and SAM could be leveraged to support biometric authentication into additional applications as well. However, this initiative could and will likely occur over time and rely heavily on Application teams and Change Control.

#### 3.8.7 Intrusion Prevention System (IPS)

Migration from IBM Proventia Intrusion Prevention to Check Point IPS system will involve Firewall and Network teams that currently operate Proventia. Additional Check Point consulting services will be required for this initiative to correctly install, configure and migrate any custom detection rules currently in operating in Proventia. Change Control and Firewall project completion dependencies must be completed prior to starting this project.

#### 3.8.8 Identity and Access Management

IBM Security Identity Manager and Security Access Manager Solutions depend on Active Directory and Multi-factor authentication with biometrics project completion. Change Control also needs to assist coordination between IBM SAM consultants and application teams for systems and applications that leverage SAM's access management capabilities. Core systems such as Email and Web Gateway authorization could me added to SAM upon completion. However, core systems such as Active Directory and biometric interfaces need to be available prior to SIM and SAM deployment.

#### 3.8.9 Email Gateway

No changes to Email Gateway architecture are proposed. Integration to McAfee ePolicy Orchestrator can take place independent of any other proposed security solution initiative. SIEM integration would need to wait until Qradar has been deemed operational. Only then will the connectors work for log collection.

#### 3.8.10 Internet Content Filter

Internet Content Filter replacement will necessitate a major change control records that will require coordination and approval across DAS/BEST and a number of agencies. Firewall and Change Control functions will need to be fully functional and prepared for the change. Local policy administrative delegation requires group member definitions and assignments associated to their respective agency departments. This effort can be planned and coordinated well before any technology solution gets deployed. Content reporting also needs to be carefully isolated per group role and delegated to each department. Written policies and standards need to support these delegated tasks and will have to be passed by senior management at DAS/BEST along with Agency leaders.

#### 3.8.11 System incident and Event Manager (SIEM)

Qradar collects various log information from security, application and other infrastructure systems throughout DAS/BEST. Timing a consulting services engagement for Qradar after most of the systems that would feed raw log data have been brought to at least test or production operation. The Qradar application can be installed and configured with a base syslog import but more valuable information and alerts will be made available after a qualified consultant fully configures and tunes the system. No other

dependency exists prior to installation but the data feeds from other systems will need to occur after they become available.

#### 3.8.12 Network device security

Network devices such as switches, routers and firewalls leverage TACACS+ for centralized authentication. DAS/BEST currently utilizes both local and centralized management accounts with TACACS+. The move to authenticating to centralized Active Directory accounts or even multi-factor tokens or biometrics will require integration into IBM SIM and SAM along with Active Directory support. Change control will also be required. Coordination with the Network Team and related support departments must be approved prior to any account changes, as this would be high service impact risk due to failure.

#### 3.8.13 VPN Remote Access

Migration to Check Point VPN remote access depends on Firewall and Change Control functions complete and operational. Check Point consulting services can create the configuration placeholders for VPN remote access and allow DAS/BEST firewall teams to slowly migrate the user connections from Cisco and Nortel. Dependencies along with the firewall configuration include Change Control, Desktop support coordination and network team for Cisco and Nortel decommissioning.

#### 3.8.14 Anti-virus, Anti-malware, Hard Drive encryption

Integrating Anti-virus, Anti-malware and Safeboot log information into SIEM depends on Qradar deployment and configuration of the data feed connectors from McAfee ePolicy Orchestrator. Information from EPO will create a much more context rich troubleshooting and Incident analysis environment for greater operational efficiency.

#### 3.8.15 Application Security

No operational changed need to take place. However, any new application changes would benefit from streamlined Change Control processes. The follow diagram shows Security Control initiatives and their dependencies. Each control in green can be started and completed independent from the other lines as long as the dependencies have been completed to the point where production integration is possible.

| Key | |
|---|---|
| Dependent system | |
| Dependent system partial component | |
| New system initiative | |

| Control Initiatives | Multi-Factor Authentication | Change Control | Firewall | Directory Services | Identity & Access Management | SIEM |
|---|---|---|---|---|---|---|
| | Dependencies | | | | | |
| Physical Data Center | | | | | | |
| Change Control | | | | | | |
| Security Configuration Management | | | | | | |
| Firewall | | | | | | |
| Directory Services | | | | | | |
| Multi-Factor Authentication | | | | | | |
| Intrusion Prevention | | | | | | |
| Identity & Access Management | | | | | | |
| Email Gateway | | | | | | |
| Internet Content Filter | | | | | | |
| SIEM | | | | | | |
| Network Device Control | | | | | | |
| VPN Remote Access | | | | | | |
| Anti-Malware and HD Encryption | | | | | | |
| Application Security | | | | | | |

Figure 33 - Security Control Initiatives and Dependency

Strategy around setting initiative priority should be to follow the order in which is displayed above. Factoring dependencies along with the highest risk reduction and operational efficiency derived this order. Actual order of implementation will also depend on professional consulting resource availability and budgetary constraints. Also noted fact is most of the existing security controls have already been purchased from very high quality vendors which leaves most of the effort around correctly integrating the solutions to gain the most value from past technology spend.

# 4 Implementation Strategy

The implementation of this strategy has four key areas. The first of these is the actual infrastructure. Implementing a solid foundation for the environment is extremely important to the overall success of the success of the overall strategy.

The second important step for the successful implementation of the strategy is a billing system for determining customer usage of the internal infrastructure at the DAS/BEST data centers. Although DAS/BEST may not actually bill customers for their usage of the infrastructure, understanding the actual costs for that infrastructure by customer is crucial to reigning in virtual sprawl – a condition of virtual infrastructures that seems to encourage the use of the environment for less critical computing.

The last two steps are closely related. These are the implementation of automation and configuration management within the overall infrastructure. These steps are needed to reduce repetitive and manual

tasks executed by the DAS/BEST IT staff. Ideally, automation would help reduce the need for manual configuration of the hardware components, which include the network, storage and servers. Configuration management addresses the automation of configuration of the operating environment and supporting application software within the overall infrastructure.

These last two steps together significantly reduce the overall effort required by the IT staff to engage in daily repetitive tasks and serve to prepare the organization for building and supporting converged infrastructure.

## 4.1 Infrastructure

The infrastructure implementation requires several parallel efforts. The first of these is the design, purchase and build of either new or additional virtual (x86) infrastructure at both of the data centers that will be in use by DAS/BEST. The second is the isolation of the remaining two infrastructure types (virtual z/VM and physical x86) so that each infrastructure has dedicated storage and the appropriate network switching in place to support access by customers. The other effort that would be integrated into the other efforts is the design, purchase and build-out of the necessary network infrastructure at both of the data centers (Groton, CT and Springfield, MA) to support the connections required by agency customers for all three infrastructures.

To complete these efforts many tasks need to be accomplished with close integration and orchestration of several teams. For the initial build/migration of the virtual (x86) infrastructure, the following steps need to be completed:

1. Determine what components in current virtual infrastructure will be kept - specifically the servers (storage will almost certainly be replaced). This would be the x86 virtual hosts currently in use that could be used for second data center site with new infrastructure in use at primary site.
2. Determine growth in virtual infrastructure needed for next two years. (This would be either a single complete cabinet or 50% of capacity).
3. Determine storage requirements for replacing current storage capacity and capacity needed for next two years of growth.
4. Determine network switch requirements (leaf/spine design using Brocade VDX switches) for virtual infrastructure rollout.
5. Complete initial detail design for virtual infrastructure to include storage, servers and network components and create Bill of Materials for quoting and ordering.
6. Purchase components needed for both sites to support 75%[17] redundant infrastructure or failover of 75% of the application infrastructure.
7. Purchase storage capacity needed at both sites – must arrive for implementation at same time as the virtual server infrastructure.
8. Purchase network switch components needed for both sites to support all virtual infrastructure to be installed (Brocade VDX switch series).

For the virtual (z/VM) mainframe virtual infrastructure, the following steps will need to be completed:

[17] Assuming that not all customers will require or need a second site for failover and recovery – otherwise this might be need to be 100% if all customers require a second site for failover.

1. Determine what components in current z/VM infrastructure will be kept together as standalone infrastructure (includes z/VM Mainframe, Brocade FC switches, and Storage Array - IBM XIV).
2. Determine network requirements to support connection of environment to core network[18].
3. Complete initial detail design for virtual z/VM infrastructure to include entire environment out to access switch layer and create Bill of Materials for the quoting and ordering of required components. This may include purchases for secondary site to support failover and recovery of applications at secondary site.
4. Complete migration plan for isolation of virtual z/VM infrastructure during move to new primary site.

For the physical (x86) infrastructure, the following steps will need to be completed:

1. Determine what components in current physical environment will be kept (specifically the servers) and which application operating environments will be migrated to virtual (x86) infrastructure.
2. Determine network requirements to support connection of environment to core network – this will almost certainly be several Brocade VDX 6710 1 GbE network switches within each rack supporting the physical server infrastructure as needed. Current Cisco Catalyst switches could also be reused to keep costs down.
3. Complete initial detail design for physical (x86) infrastructure to include entire environment out to access switch layer – this will almost certainly include network and storage. Physical server components may not change.
4. Create Bill of Materials for the quoting and ordering of required components. This may include purchases for secondary site to support failover and recovery of applications at secondary site. Network components can probably be reused; storage will need to be purchased to replace legacy hardware that is no longer under support by vendor. Additional physical servers may need to be purchased for secondary site.
5. Complete migration plan for isolation of physical infrastructure during move to primary site.

Overall, these three efforts represent the major parallel workflows that need to be completed to design, purchase and build the necessary infrastructure at both sites (primary and secondary) and migrate the current applications in production to these environments. The network design, purchase and build out must become an integral part of the build and migration of applications to these three infrastructures and will certainly require some changes to the overall network configuration and logical topology.

Once the migration to the new infrastructures at the new data center sites are complete, an ongoing effort should be maintained to continually migrate any physical applications to be hosted within the virtual infrastructure, assuming it is both technically and legally possible and does not serve to severely increase the cost to support that particular application.

## 4.2 Billing

The next major area requiring implementation is a mechanism for determining ongoing costs to support the agency customers within the three types of infrastructures. Unfortunately, the current software applications that allow for this type of billing are based on virtual infrastructure – use for physical or

[18] Some changes will almost certainly need to be made with regard to VLAN assignments, IP address assignments and routing to agency customer networks over MPLS.

virtual (z/VM) infrastructures may not be possible. Depending on the choice of billing software applications chosen, the various costs should be entered based on staffing and purchases made during the prior implementation phase for infrastructure.

This phase is critical to the ongoing success of the DAS/BEST IT organization. Since the organization is responsible for supporting multiple agencies, the need for understanding the costs incurred by supporting each individual customer and the impact to the organization budget is extremely important. Over time as customer needs increase, the general inclination of most customers, when cost is not an issue, is to make use of all available resources without regard to other customers or the agency supporting them. This would be particularly true in the case of virtual sprawl, the condition where virtual guests proliferate without bound since no cost is incurred.

Fundamentally, this is the tragedy of the commons[19]. To encourage the self-regulation of the environment by the agency customers and their use for their applications that reside within this infrastructure, some mechanism must exist to encourage it. A method for billing these customers, even if no payment is received, would help in determining which customer is making the most use of the infrastructure. This information can then be used to understand what growth will be required and which customers require these resources. It can also be used to determine polices for the assignment of additional resources to customers. This information can also be shared with all of the agency customers as a method to show the relative usage of each agency relative to the others.

This should be started immediately following the completion of the various infrastructures at both data center sites. As soon as the applications are migrated to the various physical and virtual infrastructures at the new sites, a project should be undertaken to identify the best candidate for software to be used in billing customers. Next, the choice for software should be implemented so that monthly costs are applied to each customer's account and billing reports can be disseminated among the DAS/BEST leadership team for further action as needed.

Table 15 provides a sample list of possible requirements for a financial management system (billing) that could provide useful information for planning, budgeting, billing, tracking and forecasting of the overall IT infrastructure. This is not intended as a complete list, but does provide some suggested requirements that are generally important to hosting service providers. A brief review of the goals for this software should be examined by the organization to produce an initial, though not comprehensive list, of potential useful requirements. This may occur at the same time as a review of the various software products currently available.

[19] An economic theory by Garrett Hardin – "it is the depletion of a shared resource by individuals, acting independently and rationally according to each one's self-interest, that act contrary to the group's long-term best interests by depleting the common resource" (Tragedy of the Commons) – in this case the overall virtual infrastructure.

| Requirement | Based on | Purpose |
|---|---|---|
| Capture Hardware Costs | Storage, Server, Network purchase/support costs | Needed to determine final unit costs monthly for proper billing |
| Unit Costs | Virtual Guests, Support | Needed to determine final monthly billing for customer based on unit usage |
| Planning/Budgeting | Consolidate budgets for planned/actual spending | Needed to provide variances and track accumulated savings |
| Customer Analysis | Determine how IT customers are consuming services | Needed to track costs over time and usage for various services – assists in planning for future staffing and infrastructure purchases |
| Reporting | Dashboards for resources and services for excessive cost | Needed for real time adjustments to IT organization |
| Profit and Loss | By service and customer | Needed for tracking overall health of business |
| Consumption Forecasting | Infrastructure Usage due to demand | Needed for planning purchases and budgets |

**Table 15 - Financial System (Billing) Requirements**

The software chosen should be implemented in phases. These phases are directly related to the order of implementation for the major services to be implemented – specifically the Infrastructure (IaaS) and Platform (PaaS) services defined in the services strategy. Therefore, once the software has been chosen, after the infrastructure is in production at both sites, the effort to determine costs and define billing for the infrastructure should be completed, while the effort to implement the Platform services continues.

## 4.3 Automation

Automation of the hardware infrastructure will be a key component to reduce costs within the overall infrastructure. However, to determine if actual efforts are lowering costs, some method for determining the initial costs and the costs after automation is critical. Therefore, the billing system should be implemented ahead of the automation. This allows DAS/BEST to review the costs for supporting the infrastructures over time as it relates to the billing of their customers.

Automation in this case, is the use of software tools such as VMware vOrchestrator[20], to create workflows that eliminates the need for manual effort from an administrator and reduces the execution of many individual discrete steps into a single workflow executed from within the vOrchestrator interface. The goal of automation is the elimination of all unnecessary hardware and software configuration as it relates to the creation of individual virtual guests (provisioning) for agency customers.

Ultimately, the goal is self-service or the ability of agency customers to either provision their own environment or have it done within a very short time period using automation so that they have immediate access to a platform on which the development of an application can be started without waiting for a series of discrete steps to be performed. This has the dual effect of both reducing labor costs to provide services to customers and providing a faster response to customers requesting that service – in this case the basic provisioning of an application environment.

Once the infrastructure is built and billing practices and polices implemented, automation of those tasks that are completed most often within the infrastructure should begin. The approach should be to

[20] This software (application) is available at no extra cost with the implementation of vCenter Server within the virtual infrastructure.

eliminate the manual execution of groups of related repetitive tasks and have these tasks executed as a single automated workflow without the need for manual effort or intervention.

The majority of these tasks should be prioritized based on the frequency of occurrence. For example, if the task of making firewall changes occurs on a weekly basis, but the task to provision a firewall occurs quarterly; then the changes to the firewall should be automated first. If the virtual guests running in a production environment need periodic configuration and it requires 5% of the administrators time, but provisioning requires 15% of the administrators time, then the provisioning tasks should be automated first. Once the top 20% of the prioritized tasks are automated to a sufficient extent, the next 20% of the tasks should be automated.

This iterative process should be continued until the only tasks left to automate are those that happen so infrequently that automation may not be cost-effective - automate tasks until the only tasks left happen so infrequently that automation doesn't make sense.

Automation of the infrastructure should be pursued within the virtual infrastructure using standard interfaces such as RestFul Web services for all of the major components – servers, storage, and networks. For the remaining two infrastructures, virtual z/VM and physical (x86), the choice to automate tasks associated with supporting those environments should be based on the expected future need for support. If these environments will slowly be migrated to the virtual (x86) infrastructure, then it may not make sense to automate these infrastructures, if, however, DAS/BEST intends to support these for their customers and their customers are expecting to need them into the foreseeable future, then it may be prudent to automate tasks within these other infrastructures as well. Table 16 lists some possible workflows that should be automated first.

| Step | Automation | Components/Workflow | Comments |
|---|---|---|---|
| 1. | Virtual Guest Provisioning | VMware: Creation of a virtual guest within the infrastructure based on specific requirements from the customer | Probably the most obvious first workflow to create, might need to be expanded to include an entire application architecture that is composed of several software components |
| 2. | Increase Capacity | Increase capacity within infrastructure. Bring new virtual host online or add additional storage capacity | This is the second most obvious choice to automate. Storage or virtual hosts would be added to the overall infrastructure, specific to a cluster. |
| 3. | Virtual Infrastructure Changes | Modify various aspects of the virtual infrastructure within a specific cluster | Examples include virtual distributed switches, datastores, security settings, and role based access. This would be executed based on customer requests |
| 4. | Reporting | Various reports on the overall infrastructure | This might include a report on the virtual sprawl by customer. |

Table 16 - Automation Implementation Schedule

These certainly are not all of the possible workflows that could be automated, but do represent some obvious choices and provide an example of the type of automation that could and should be undertaken.

## 4.4 Configuration Management

Configuration management should be the final phase of the overall implementation of converged infrastructure. This phase addresses the need to reduce the effort to build, support and maintain the operating environments within the virtual infrastructure. These operating environments run the actual application components within them and include Microsoft IIS, .NET and SQL Server. Other examples include those components supported by Linux operating environments.

Although this is assigned to the last phase, this effort could be started shortly after the automation efforts are started since two entirely different groups – infrastructure services for automation and application services for configuration management, would undertake these two parallel efforts.

For the same reason as given for the automation effort, the configuration management efforts should be started after a proper billing or financial management system is in place to track the overall infrastructure costs. This should include capturing the costs for any efforts by staff to support and troubleshoot the infrastructure and operating environments in use by agency customers – the application services being provided directly to the customers.

The actual rollout of the configuration management system should follow a bottom up approach where the lowest level and therefore most common configuration tasks are handled first with each successive higher layer handled with configuration management tools. The overall effort should slowly cease as it becomes clear that further effort will yield only modest returns to saving costs – the law of diminishing returns.

Over time as the actual application services are refined and simplified[21], the implementation of a configuration management system should require less effort to maintain and support. Further efforts by application services to define various types of application frameworks would help considerably in the implementation of configuration management by simplifying the overall application architecture that needs to be maintained by the configuration management system. Table 17 lists the various major steps within the phase for implementation.

| Step | Managed Layer | Tasks | Comments |
|---|---|---|---|
| 1. | Windows Operating Environment | Patches, updates, and configuration changes | This would be beyond the normal windows updates or patches provided by a WSUS server |
| 2. | Linux Operating Environment | Patches, updates, and configuration changes | This would be beyond the normal windows updates or patches provided by a update server |
| 3a. | Windows IIS | Installation, configuration, and changes of components | Might include the initial creation of the web server and configuration for a particular application |
| 3b. | Windows .NET | | Installation and configuration |
| 3c. | Windows SQL Server | | Initial creation of a database and tables |
| 4a. | Linux-based Web Server | Installation, configuration, and changes of components | Might include the initial creation of the web server and configuration for a particular application |
| 4b. | Java/Middleware | | Installation and configuration |
| 4c. | Oracle/IBM DB2 | | Initial creation of a database and tables |
| 5. | Code Releases | Installation, Updates | Primarily for the release of new application changes during scheduled change windows |
| 6. | Firewall/Security | Configuration and Updates | Assuming virtual network features are provided by the infrastructure (firewall, NAT, router, SSL-VPN, etc.), this would provide a completely configured virtual appliance[22] to support these features |

**Table 17 - Configuration Management Implementation Schedule**

[21] This refers to a measurement based on the actual number of different components supported from various vendors such as IBM, Microsoft, and Oracle.

[22] Several virtual appliances exist such as the Brocade Vyatta Router and SonicWall Virtual appliance to providing these features. VMware will release an appliance to support these features within the next year.

# Appendix A – TCO Analysis

StrategyModel-TCO v1-4.xlsx